NAVAL
POSTGRADUATE
SCHOOL

MONTEREY, CALIFORNIA

THESIS

SYSTEM EVALUATION OF HARDWARE AND SOFTWARE FOR A STREAMING MULTIMEDIA SERVER USING THE MULTICASTING PROTOCOL

by

John Carls

September 2003

Thesis Advisor: Geoffrey Xie
Second Advisor: John Gibson

Approved for public release; distribution is unlimited.

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE		*Form Approved OMB No. 0704-0188*
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.		
1. AGENCY USE ONLY *(Leave blank)*	**2. REPORT DATE** September 2003	**3. REPORT TYPE AND DATES COVERED** Master's Thesis
4. TITLE AND SUBTITLE: System Evaluation of Hardware and Software for a Streaming Multimedia Server Using the Multicasting Protocol		**5. FUNDING NUMBERS**
6. AUTHOR(S) John Carls		

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
    Naval Postgraduate School
    Monterey, CA 93943-5000
		8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A		**10. SPONSORING/MONITORING AGENCY REPORT NUMBER**
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.		
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.		**12b. DISTRIBUTION CODE** A

13. ABSTRACT *(maximum 200 words)*

The next step in the evolution of services provided on an intranet or the internet will be distributed or distance learning with collaboration among peers. Currently, this is done on a one-to-one basis. To expand to a one-to-many collaboration environment, there needs to be a server distributing the multimedia content without creating additional network traffic even though many users are accessing or viewing the multimedia content. Multicasting allows many users to view multimedia content without creating additional network traffic. There is a server providing multimedia content to a multicast address so users may access it. This thesis defines metrics and conducts a comparison of different servers capable of distributing multimedia content using the multicasting protocol.

14. SUBJECT TERMS Multicast, Server, Network, Multimedia, Streaming			**15. NUMBER OF PAGES** 79
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	**18. SECURITY CLASSIFICATION OF THIS PAGE** Unclassified	**19. SECURITY CLASSIFICATION OF ABSTRACT** Unclassified	**20. LIMITATION OF ABSTRACT** UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18

THIS PAGE INTENTIONALLY LEFT BLANK

SYSTEM EVALUATION OF HARDWARE AND SOFTWARE FOR A STREAMING MULTIMEDIA SERVER USING THE MULTICASTING PROTOCOL

John W. Carls
Lieutenant, United States Navy
B.S., Old Dominion University, 1998

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

from the

NAVAL POSTGRADUATE SCHOOL
September 2003

Author: John W. Carls

Approved by: Geoffrey Xie
                    Thesis Advisor

                    John Gibson
                    Co-Advisor

                    Peter Denning
                    Chairman, Department of Computer Science

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The next step in the evolution of services provided on an intranet or the internet will be distributed or distance learning with collaboration among peers. Currently, this is done on a one-to-one basis. To expand to a one-to-many collaboration environment, there needs to be a server distributing the multimedia content without creating additional network traffic even though many users are accessing or viewing the multimedia content. Multicasting allows many users to view multimedia content without creating additional network traffic. There is a server providing multimedia content to a multicast address so users may access it. This thesis defines metrics and conducts a comparison of different servers capable of distributing multimedia content using the multicasting protocol.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

LIST OF FIGURES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

Providing streaming multimedia is the next step in the evolution of Intranet and Internet services and is taking root. The streaming multimedia service will provide distributive learning (a form of collaborative video-teleconferencing where the classroom learning environment is conducted in near real-time) (See Figure 1) and distance learning (a virtual classroom of students across the country or world where the classroom learning environment is conducted in asynchronous fashion of postal or electronic media exchange of learning material, i.e. lectures, videos, and tests) (See Figure 2). The current model of streaming multimedia over the Internet or on an intranet requires high usage of bandwidth because the streaming multimedia is an on-demand service – where the user clicks the requested content and the material is provided to the client. No multicasting is used with this method. A new stream of network data traffic must be generated for each user. This creates a significant amount of traffic on a network that may not have sufficient bandwidth. A virtual classroom across the country takes dedicated circuits of high bandwidth and is very expensive to operate. Time, effort, and money are used to establish the connection from the sending classroom to the receiving classroom; dedicating circuits throughout the duration of the class. Current collaboration between classrooms is accomplished through proprietary programs, but the programs are bandwidth intensive and limit one-to-one video-teleconference. With the use of multicasting, this can be changed to a one-to-many, many-to-many, or many-to-one video-teleconference. Classroom lectures are established without creating additional traffic on the network no matter how many students wish to view the lecture. [KO01]
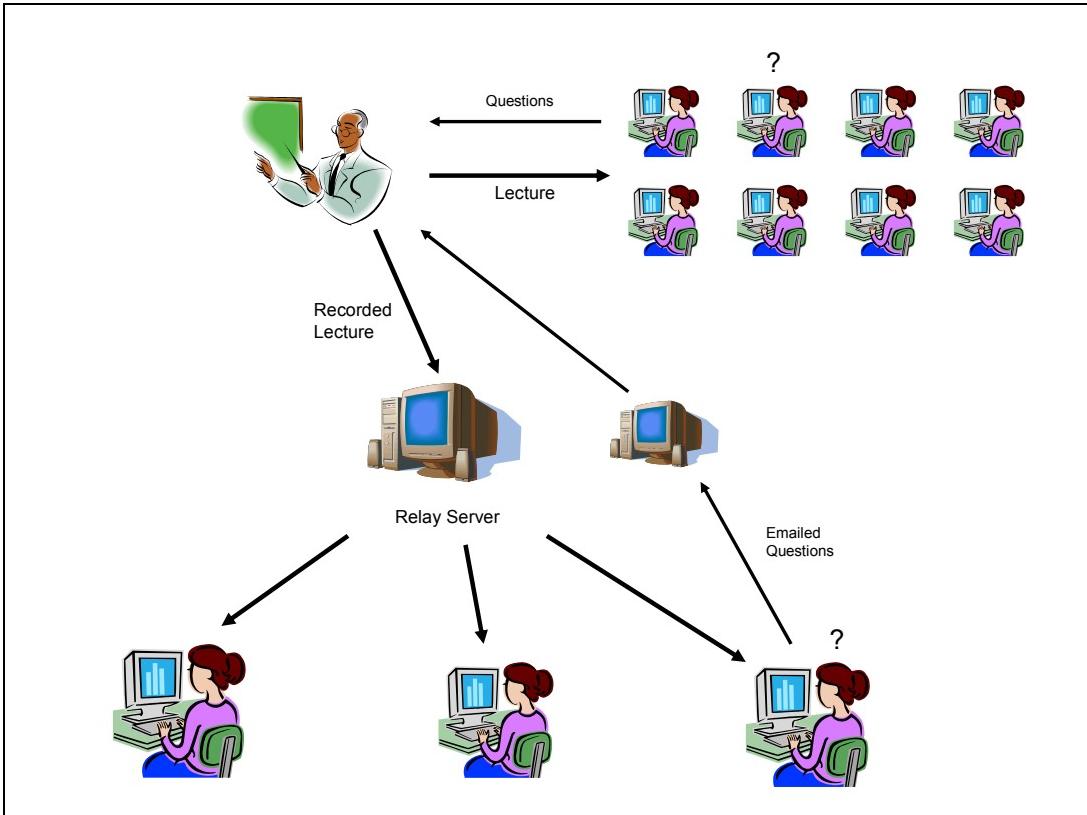
Figure 1.   Distributive Learning
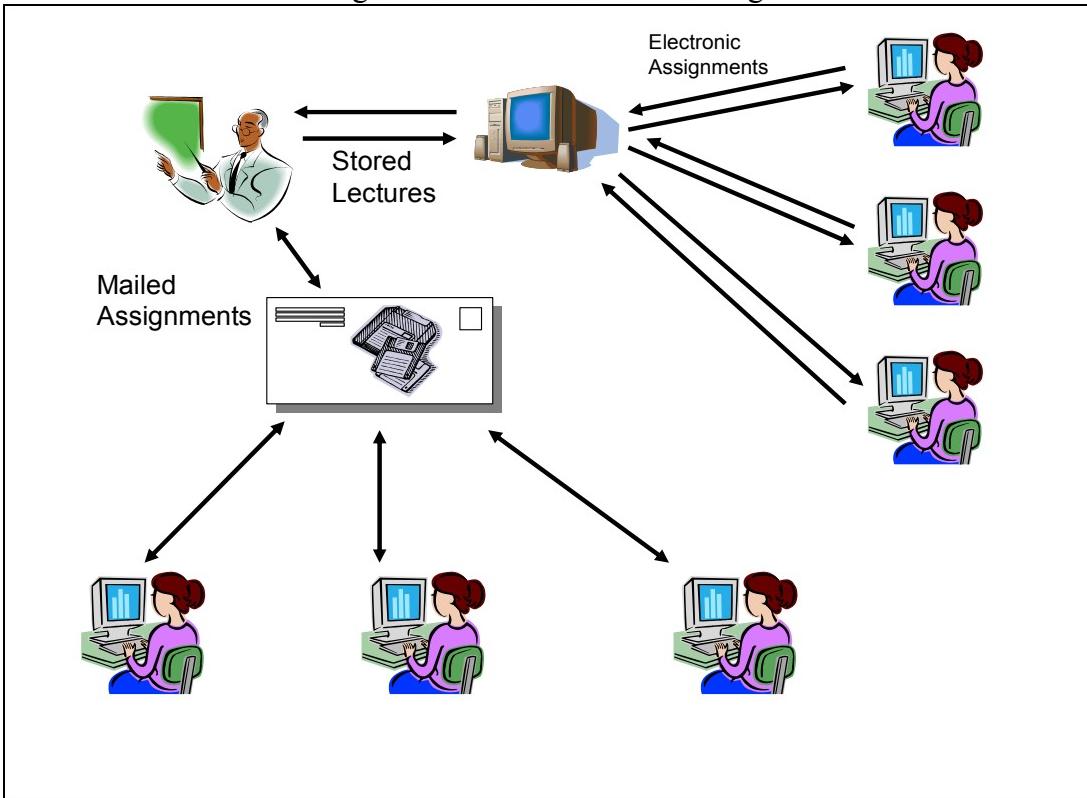
Figure 2.   Distance Learning

An alternative to the classroom lecture broadcast is allowing each user to sign in to a web page and create a new multimedia stream to their computer or workstation. This allows the student to view the classroom lecture from a remote location. Each student that requests the classroom lecture creates a new multimedia stream on the network. One of the disadvantages to this method is the one-way conversation, the instructor presents the material to the students and the students can not respond to the professor with questions unless a separate e-mail or web page is established allowing questions to be submitted. Another disadvantage is that each user creates a new multimedia stream creating an increase of bandwidth for each multimedia stream that is sent back to the instructor. This is not feasible for use over the Internet or Intranet because of the scalability required on the networks due to the amount of bandwidth that may be required to create each multimedia stream. The network may be able to handle approximately 15 students for one lecture (15 multimedia streams), but not 15 classes of 15 students each (225 multimedia streams). In the network core where traffic volume is very large, routers do not have the buffering or processing resources required to handle large multiple multimedia streams. With multicasting, this volume of traffic would be reduced to one multimedia stream per class; therefore allowing the students of the 15 classes to see the lecture on the network.

The Streaming Multimedia Multicast Server is a combination of software and hardware required to take an input from an audio and video source, convert it to digital information for streaming, and distribute over on a network to a multicast address. The multicast multimedia stream is received by the requesting computer using a multimedia player capable of receiving multicasting; such as Real Player, QuickTime, or Windows Media Player. Multicasting occurs in three methods: One-to-Many, Many-to-Many, and Many-to-One. This allows distance learning and collaboration to enhance the educational or work environment.

Thesis will provide the necessary background to understand the basics of multimedia, streaming, networking, multicasting, and system composition. By the end of the thesis, the reader should have a basic understanding of the terminology used and the metrics that defines the system comparison.

The thesis proposes to determine, specifically, the metrics for evaluation of the different multicast systems available. The research will provide the methodology used to determine which system has the best possible combination of software and hardware to create a high quality, low bandwidth streaming multimedia server using the multicast protocol. The research will entail determining the different types of multimedia hardware server systems, the different types of operating systems that are used with their respective hardware server platform, and the different streaming multimedia applications on the respective operating system and hardware server. The main focus of the thesis is on determining the metrics and the methodology used in comparing different streaming multimedia systems. The hardware and software platforms example used in the comparison will be the comparison of Windows on Intel/AMD, Linux on Intel/AMD, Mac OS X on PowerPC G3, and stand-alone equipment using the Streaming Distribution System, like the one produced by VBrick Systems, Inc.

The comparison methodology of the metrics that are determined in the thesis will be used on the different server platforms to show the comparison of the metrics used to determine which sample system provides the best cost to performance value. The results will be used as a recommendation to Code 05 IT Department as a starting point for the Distance Learning Multicast Servers. There will be recommended follow-on research relating to the thesis after the conclusion and recommendations.

II.   BACKGROUND

A.   NETWORKING

This chapter will provide the basic background necessary to understand the remainder of the thesis. The terminology will be covered to give a brief description of the various pertinent areas of multimedia, the relevant networking aspects, and the hardware and software components and systems. Multicasting will be discussed in greater detail at the end of this chapter.
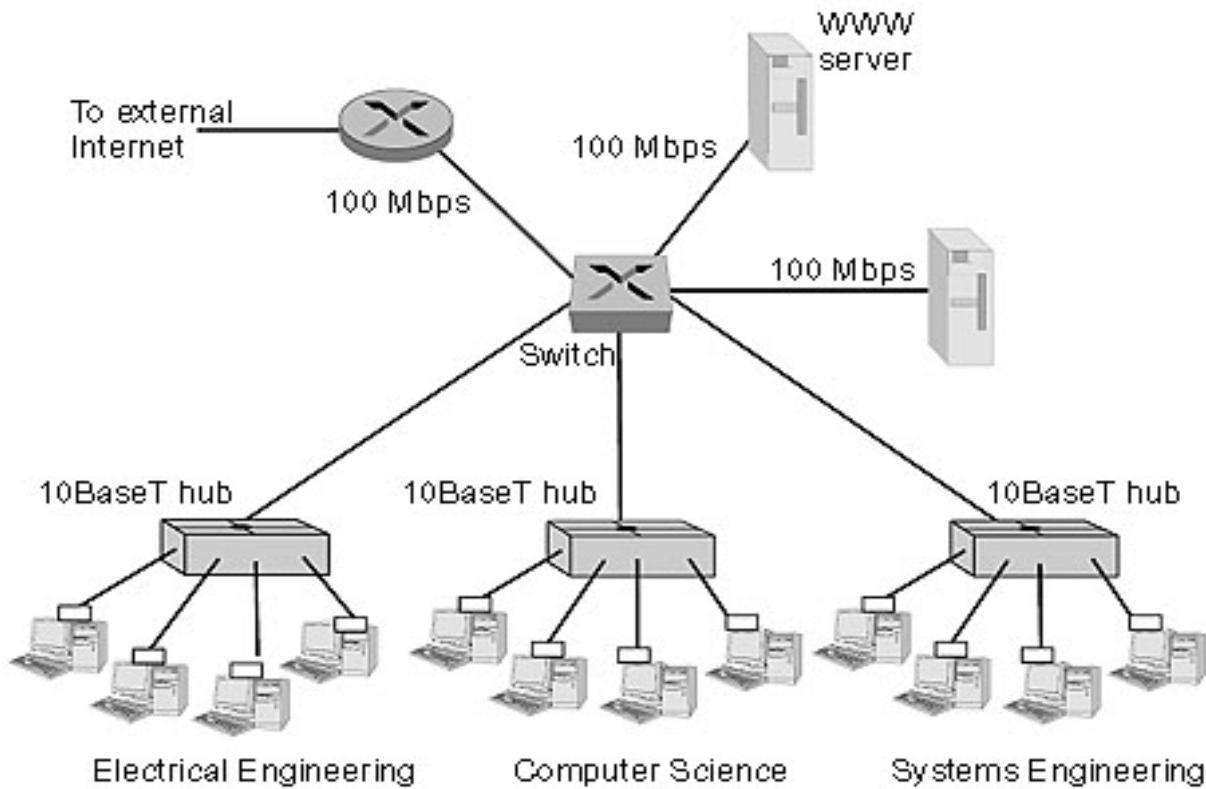
Figure 3.   Basic Network [KURO01]

The basic components required to provide a network and to share information between the users of the network are the servers, workstations, routers, switches, bridges, and connection media, as well as the protocols for the equipment to communicate with each other. There is additional equipment used in the enterprise environment that expands the basic network. The additional equipment may be a firewall, a network storage device, and a variety of specialized servers such as a proxy server, DHCP server, web server, and/or a database server. The servers provide data, pictures, music, or video

for the users on the network. The connection medium will take the raw signals from one physical point to another. In today's environment, it may be a cable, either copper or fiber optic, or wireless, either through the use of radio signals, laser, or light. The connection medium will not be covered in this study. There is a wealth of research and information available for networking. A good place to start is _Networking: The Complete Reference_ by Craig Zacker. Routers, switches, and hubs are used to assist the data bit flow from the server to the client computer. The router helps decide which interconnected networks the information needs to travel in order to get to the end user. This is done through logical addressing of the individual components, typically through the internet protocol (IP) addressing scheme, at the network layer of the TCP/IP five-layer protocol stack. The interactions of the various protocols, as depicted in Figure 4, support the exchange of data between network hosts. The last stop for the data is the client or end user. This is where the information is then displayed in the appropriate format for the user to understand. [KURO01/ZACK01]
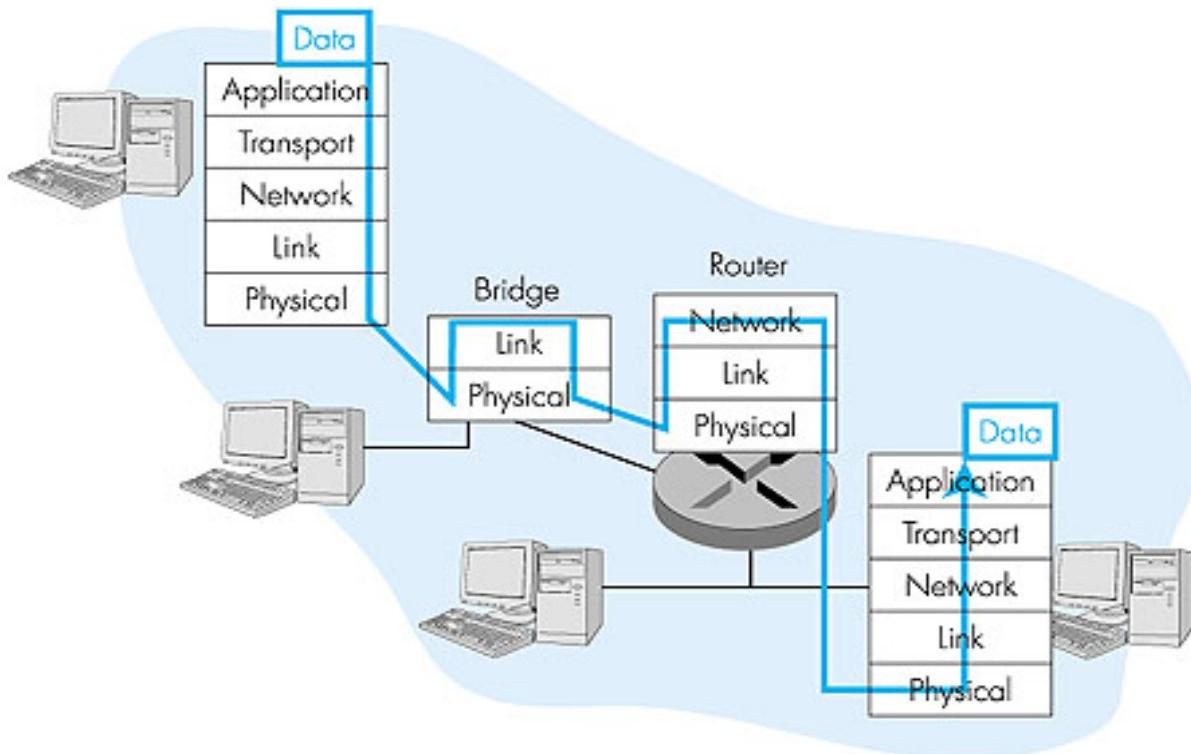
Figure 4.   Data Packet Traversal Using the TCP/IP Protocol Stack [KURO01]

1. Server

There are many different types of servers used in networking. A server is basically a computer that is dedicated to providing a particular service to other computers that request that service. The most familiar types of servers are web servers. The web server provides web pages for users looking for information on the World Wide Web or internet. Database servers are used for storage of data in some form of a record format, such as driver's license information, insurance records, and criminal records, which are readily manipulated or accessed using a manipulation language such as the Structured Query Language. An e-mail server will store, send, and retrieve electronic mail for users associated with the particular company hosting the e-mail server. The file server, or ftp server, server provides storage and sharing of various media files. Other servers include Dynamic Host Control Protocol, Network Area Storage Servers, Proxy Servers, Backup Servers, and Streaming Multimedia Servers. Only the later will be discussed in the thesis. This server provides support to distance learning by transmitting a multimedia stream to a multicast protocol address that remote users may access for viewing.

2. Routers / Switches / Bridges

The group of equipment referred to as routers, switches, bridges, and hubs is often confused or misinterpreted. Each piece of equipment performs a specific function at a different layer in the protocol stack. Some manufactures provide multiple functions or have combined the different layers and functionalities in the equipment, which can lead to confusion for the user.

The router is a networking device that connects two or more separate local or wide area networks with different addresses together. In so doing, the router isolates the connected networks into separate broadcast domains. The router works at the network layer (layer 3) of the OSI model, which allows the router to connect different underlying link protocols, such as Ethernet and Token Ring. A router will store and then forward the data packet to the appropriate network based on the network layer address. The switch is often confused with the router, but the main difference is that the router works at layer 3 while the switch works at layer 2, using physical addresses. The router has the ability to

7

handle multicasting and will enable (graft) or disable (prune) the data stream to the users as required.

The switch or bridge allows multiple connections but does not necessarily broadcast traffic on all connections, just on the one for which the data is specifically intended, assuming the switch knows to which port the destination address is connected. The switch or bridge is often used on a segment where there are multiple users connected, such as a computer lab. The switch will read the incoming packet address and then send the packet to the appropriate port according to the table that is stored in the switch (store-and-forward). When the switch is initially installed, it must learn the addresses of the devices connected. This is done by reading the source address when a packet comes in on a port from the connected device. The switch works at the data link layer (layer 2) of the OSI model, which allows the switch to connect similar segments of a subnet together often being used at the end of a segment and branches to users. Some of the more sophisticated switches have the ability to forward multicasting traffic to the appropriate user as required, while others can not handle the multicast traffic and will either broadcast it to all users on the switch or delete the packet altogether. This can cause problems on the network, so it is important to take note on the capabilities of the switch being used. Note that while routers establish distinct broadcast domains, switches and bridges establish separate collision domains for contention-based protocols or separate rings for token-based networks.

The hubs and repeaters all work at the physical layer (layer 1) of the OSI model. These devices are basically designed to amplify the signal and pass it along allowing the network length to be extended. They do not provide any data filtering, addressing, or segmentation of the network traffic. The hubs and repeaters do not have any capability to process multicasting traffic. The device will just pass on the packet of data as if it where a unicast or broadcast data packet. Essentially, a hub is a multi-port repeater.

Routers, switches, bridges, hubs, and communication medium provide the basic communications infrastructure for the network. The important thing to remember about these devices is their functionality and what layer of the OSI model they function. For the thesis, it is important to make sure that the devices connected to the network have the

capability to handle multicasting. The servers tested use multicasting and therefore, must be appropriately supported by the other network devices to function properly. Without the ability to handle multicasting, the streaming multimedia servers would not be able to function as required. Some devices are not capable of handling multicasting and can cause major problems on the network by broadcasting or flooding the packets to all users. This would defeat the purpose of having a multicasting multimedia server for distance learning, essentially functioning as a unicast multimedia server. [KURO01]

3.   Client / User

The client or user is the computer or the human requesting a service from the server, whether it is e-mail, web pages, a database record, or other data. In the case of this thesis, the service will be to view a multicast multimedia stream for distance learning.   There are many varieties of computer hardware, operating systems, and applications that are available to use and view the information available on the network. Some of the hardware is capable of using different operating systems while other hardware can only use one specific operating system. The applications can be multi-platform independent or may only run on one specific system. The client system configuration, to include hardware, operating system, and application software, is not discussed in the thesis, as it will be the up to the user to determine the hardware and software configuration of the client computer.

4.   Unicast

As noted, unicast traffic is generated when one computer sends a data packet to another computer across a network. This may be referred to as a point-to-point or a one-to-one communication. The size of the data stream created is directly proportional to the amount of data that is to be transferred. For each user wanting to receive a multimedia stream, an individual unicast stream must be transmitted. This can quickly consume the bandwidth available on a network segment. For example, if there are 50 users on one segment and each user requires a data stream greater than 2 Mbps, then a 100 Mbps segment will not be able to handle all the data required by the users. Even if the users want the same multimedia stream, the stream must be replicated 50 times. This does not include any overhead that is required by the protocols to process and route the data packets. For multipoint-to-multipoint group communications, given all the users wanted

to communicate with every other user using a unicast connection, the number of connections required would be (N^2 – N) / 2 connections, where N is the number of users. For 50 users, this would require 1,225 connections. If each connection is 2 Mbps, the total bandwidth would require over 2,450 Mbps (over 2 Gigabits per second – 2 Gbps). Many network backbones would not be able to handle this amount of traffic, and certainly not a 10 Mbps or 100 Mbps network segment. This is a potential disadvantage of using unicast for group communications. [WITT99]
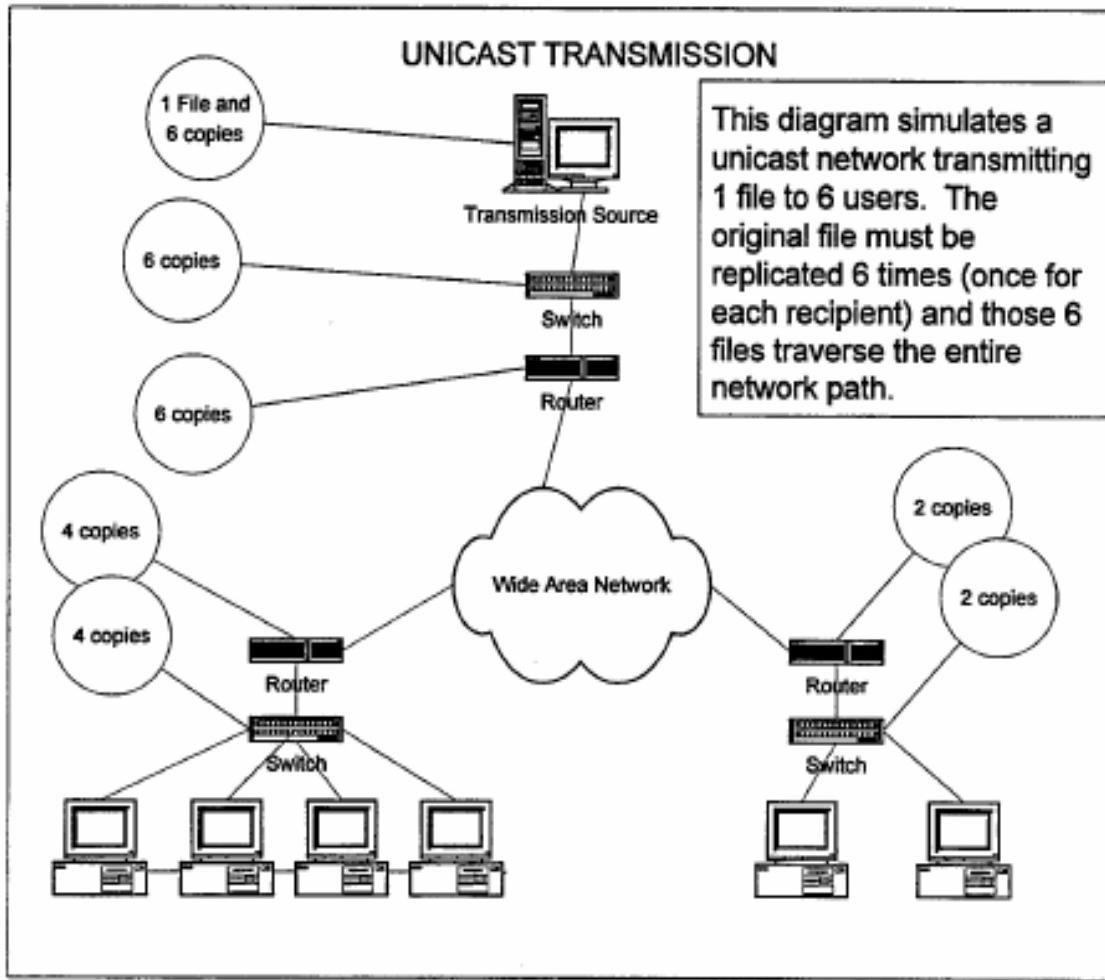
Figure 5.   Unicast [FLOU99]

5.   Broadcast

Broadcast network traffic is when one computer sends a data packet to all computers across a network. This is a one-to-many or a one-to-all connection. For each broadcast data stream sent from a server or client, there is only one individual stream. There is only one broadcast address so a distinction is not needed between the broadcast

addresses. The broadcast traffic only uses the required bandwidth for the stream and does not consume additional bandwidth on the network segment, although it requires all hosts on the segment to process the stream. If there are 50 users on one segment and the server sends a broadcast data stream that is 2 Mbps, then a 100 Mbps segment will easily handle all the data required to the users. The data stream is transmitted once and not replicated 50 times. Broadcasting is a form of group communications, and broadcast traffic is sent to all the users whether they client wanted to receive the communication or not. This is a disadvantage of using broadcast for group communications. [WITT99]

B.   STREAMING MULTIMEDIA

Streaming multimedia, or multimedia networking, comes in a variety of formats and includes terminology that is often confused. The terms used throughout the thesis are clarified below.

1.   Multimedia

Multimedia can mean several different things, depending on the application and the environment in which it is used. Multimedia is can be thought of as multiple mediums combined together and usually consists of video and audio. Sometimes audio alone is considered multimedia even though it is only one medium. For the thesis, multimedia includes the use of audio, or sound, and video in an application program.

In art, multimedia refers to the different types of material and media used to create a sculpture, painting, or presentation display. In a teaching or learning environment, multimedia can refer to the different types of media available such as overhead projectors, pictures, video tapes, or audio tapes. These items may still be used in distributed or distance learning; however, they are used sparingly because the material may be difficult to portray across a small video stream or through the mail. When computers and the internet were in their early development stages, the technology did not have the capability to play both video and audio together. The data files were often large, several Megabytes (MB), and the connections were often slow, 14.4 kilobits per second (kbps) which resulted in a long wait to listen to music or watch a video clip. The technology available today has improved tremendously and it is possible to watch 30

frames per second, full-motion, 640x480 full-screen video, with CD stereo quality audio. Standards have evolved for audio and video. [MICR03]

a.   Audio

The following is the basic breakdown of audio types and the different formats available. Audio types can be differentiated by the number of channels used to provide the sound quality. The first and lowest quality level is monophonic. There is only one channel available for sound or audio. The next level is stereophonic. This provides a left and a right channel for sound. The newer formats provide a surround sound with multiple channels for sound, such as Dolby Surround Sound 5.1, using 6 channels. Most computer technology today does not have this available for recording or playback, unless the system has specialized components. The thesis is only concerned with the monophonic or stereophonic formats.

The next consideration for the audio is the recording format. When recorded audio was introduced, the recording format was accomplished in an analog fashion. Phonographs and tape decks are analog; the equipment has motorized moving parts that make the recordings.   Today, audio is recorded digitally, where the equipment making the recordings is made of discrete electrical components. To accomplish a digital recording, the audio signal is sampled and the discrete samples are quantized, resulting in a finite set of sample levels. The quality of the playback depends on number of bits allocated for each digitized audio sample.

Initial computer and digital recordings used only 8 bits per sample, resulting in 256 different values, and a low sampling rate. The standard used for compact disc quality recording and playback, however, uses sixteen bits per sample, yielding 65,536 different values. Increasing the bits used per sample to 24 bits will result in over 16 million values. The more bits used, the more accurate will be the sample recorded. The sampling rate determines how many samples are taken in a given amount of time, measured in samples per second.   The CD Stereo-quality sampling rate is 44,100 samples per second. With 16 bits per sample, this yields a data rate of 1,411,200 bits per second, or approximately 1,378 Kbps. The sampling rate will make a big difference in the size of the audio stream that will be sent across the network. Note that sampling rates

are often reported in hertz, rather than samples per second. One hertz is one sample per second. Common sampling rates used with computers over the last 10 years include 4 kHz (mono), 8 kHz (telephone mono), 16 kHz (G.722 compression standard), 22 kHz (½ CD stereo rate), 44.1 kHz (CD stereo), and 49 kHz (DAT stereo). [MACK02]

b.   *Video*

The second part of the multimedia combination is the video.   As mentioned earlier, in its infancy, the Internet could not handle large data files. Now that technology has significantly improved, there is sufficient bandwidth available on most networks to carry both audio and video data. Video resolution has improved, signaling has changed to digital, and color has been added.   The original television screen resolution was 720 horizontal lines and 540 vertical lines, giving an aspect ratio of 4:3. However, not all the lines are viewable, as some are used for timing of the refresh rate for the pixels.

Each line is composed of individual picture elements (pixel). The visible resolution for television is 640x480 pixels.   High Definition Television (HDTV) has increased the aspect ratio to 16:9 resulting in two screen pixel resolutions of 1920x1080 or 1280x720 using high resolution.   A lower standard resolution of the HDTV 16:9 aspect ratio displays 704x480 pixels. In addition to the number pixels that are used for the screen resolution, there is also the rate at which the pixels are refreshed. For the United States, most of North America, Japan, and most of South East Asia, the refresh rate is 60 Hertz. For most European countries the refresh rate is 50 Hz. Both of these standards are based on the AC rate of the regions' power systems. This is the major reason why there is a difference in the technology used in the United States versus the technology used in Europe. The horizontal lines are divided into two groups, called frames, with each group skipping every other line. Each frame is refreshed 30 times per second, with the frames being interleaved. Although the frame rate for recoding has stayed the same, improvements in technology do allow for a higher frame rate. The human eye can not detect the difference for anything higher than 24 to 30 frames per second, so applications do not need to be any faster than 30 frames per second. Settings used in video applications range from 30 frames per second to 1 frame per second, were the higher the frame rate, the better the quality. The frame rate should be set according to

the needs of the supported application to ensure only the amount of data required is sent across the network. For action video, the frame rate should be set at 30 frames per second, where a low-action topic, such as a person who is just talking, could be set as low as 10 frames per second and still provide a quality presentation. In addition to the frame rate and pixel resolution, the depth of color must be considered. Using 8 bits for depth for color only results in 256 color variations, where as a 24 bit-depth will yield over 16 million colors. All this information must be encoded at one end and then decoded on the user's end, so this must be taken into account. A small video clip that is recorded for 5 seconds at the standard television rate of 30 frames per second, a screen resolution of 640x480 pixels and a 256-color pallet will result in a file size of nearly 44 megabytes. One black and white still-frame from a 640x480 video clip would be 37.5 kilobytes. Add 256 color resolution (8 bit color depth) and the file increases to 300 kilobytes. One second of motion results in a file that is almost 9 MB. But reducing the frame rate to 15 and a screen resolution of 320x240 and 16 colors will significantly reduce the same five second video clip file to under 3 MB. While more tolerable, a bandwidth constrained network would have difficulty supporting such a file transfer. This data-rate does not even include the associated audio data! [BRAI31/BRAI32/BROW03]

c. Standards

During the 1980s, when it was realized that telecommunications and the computer industry were becoming more related and intertwined, the International Telecommunication Union (ITU) developed the H.261 and H.263 video conferencing standard that is the basic foundation of today's multimedia standards. The standards were designed to use a single stream with timing information, which is an important function. Once the audio and video are combined into a single stream, the data are in a form well suited to digital storage or transmission. The principal format in use today is by the Moving Pictures Experts Group. The group was formed in 1988 to devise a system in which encoded video could be used by applications and stored on CD-ROM. The first standard, developed in 1992, was the MPEG-1 format and it was designed to operate with a maximum transfer rate of 1.5 Mbps, providing a quality level equivalent to VHS. In 1994, MPEG-2 was released to incorporate higher transfer rates, larger pictures, and use with DVD devices. MPEG-4 became the standard for the Internet in 1999. It

provides generic descriptors to the Quality of Service (QoS) standards available on modern networks. For more detailed information about the MPEG standards available, visit the web site for the Motion Picture Experts Group at http://www.chiariglione.org/mpeg/index.htm. In addition to the MPEG standards, there are a variety of proprietary standards available by various corporations.

2. Download vs. Streaming

There are two ways of viewing multimedia files. The first is to download the file from the server to the client computer and then open the file and play the multimedia clip. The second way is to start the file transfer from the server to the client computer and once a small portion is buffered on the client computer, the multimedia clip starts playing. The latter is what is known as streaming. The problem with downloading is that once the file is viewed, it may no longer be required. If so, the user must remember to delete the file when finished. If the file is not deleted and other files are downloaded, eventually the client hard drive will run out of disk space. This could leads to additional work for the IT Department because the technicians must come in and delete the unnecessary files if the individual that downloaded them is no longer available. If the typical multimedia download file is 500 MB in size, it won't take long to fill a 10 GB hard drive. Storing the file to a mapped directory also risks filling limited storage space again leading to storage problems for the IT Department. Additionally, if multiple personnel download the same file multiple copies of the file may reside on various client computers or mapped directories, clogging the shared resources. With streaming, the file is stored in one location on a server, can be played by hundreds of users across the network, and there is no additional storage required because once the multimedia file is played on the client computer, the file is deleted. [HOST03]

C.   SYSTEM COMPONENTS

The Streaming Multimedia Server is comprised of various individual hardware and software components. The server system will need an operating system to control the functions and operations of the system. As certain operating systems can only operate on certain hardware platforms, the hardware will usually dictate what operating system is being used for the multimedia server. Also required is the streaming multimedia application that performs the functions necessary to provide the multimedia streaming

services to the client user. The applications must be able to support multicasting and streaming multimedia concurrently.

1.   Hardware

The hardware for the Streaming Multimedia Server has two basic forms. The first form is the stand-alone unit. The stand-alone unit consists of the hardware to perform the appropriate functions in a self-contained fixture. There is no user capability to upgrade to newer equipment without replacing the entire device. The functions consist of receiving audio and video data, either analog or digital, converting them to digital data that can be electronically stored and then be streamed out across the network to a multicast address and viewed by the user(s). The second form is the general purpose computer system. The computer system can be upgraded with newer or more capable hardware or software components, such as adding more memory, a higher quality audio/video encoding/decoding graphics card, more hard drive storage, improved applications or operating system software, or later device driver software. The computer system may be either open source or proprietary. The open source computer uses components that are available through various venders and may use various operating systems, such as Windows XP or Red Hat Linux. The proprietary computer can only be purchased from one vender, and if any upgrades are needed or wanted for the server, the components must be purchased from the same vender. An example of a proprietary computer system is Apple or Sun Microsystems. The general purpose computer system platform has the same functionality as the stand-alone unit but with more features and robustness.

2.   Software

The operating system for the Streaming Multimedia Server may also be either open source or proprietary. However, there are a wide variety of vender versions for both open source and proprietary operating systems that range from free and unsupported to expensive with technical support. The most popular and common operating systems available are Windows, Linux, UNIX, and MAC OS. All of the operating systems available have multiple versions on the market. The latest version generally incorporates the latest hardware capabilities and features and functions that were not released in the previous version.   The operating systems from Windows and Apple are considered proprietary while the Linux distribution from Red Hat is open source. For UNIX and

16

Linux, there are various open source venders that offer different variants of the same basic underlying architecture. The various different venders are Red Hat (Linux), Caldera (Linux), SUN Microsystems (Solaris UNIX), Hewlett-Packard (HP-UX UNIX), Berkley Systems (FreeBSD Linux), and SCO Linux. For more information about the differences between the Windows, Linux, and UNIX operating systems, see *Modern Operating Systems* by Andrew Tannenbaum.

For the stand-alone units, the proprietary operating system is streamlined and is not easily upgraded. This type of operating system is typically implemented in firmware, which is embedding the compiled code in an integrated circuit like Read Only Memory (ROM). The operating system only offers the basic functions that the system is designed to do and may allow for some limited parameter configuration. Whether the operating system is embedded, proprietary or open source, the Streaming Multimedia Server operating system must be able to support the hardware being used and streaming multimedia files to a multicast address on the network.

3.   Applications

The application for the Streaming Multimedia Server is the software program that serves the content. It is not part of the operating system. The application program is responsible for converting and encoding the audio/video signal. It interfaces with the transport and network layer software to send the data across the network to remote client computers. All the conversion and digitizing is done "behind the scenes" and transparent to the users/clients. The applications typically provide a graphical user interface (GUI) so the operator can perform the necessary functions by "pointing and clicking" the mouse. A command line interface is very cumbersome and difficult to use for the novice user, which will deter users from purchasing such applications, especially if there is one available with a more user-friendly GUI. Within the application, the multicast address for the respective stream may be specified; the audio mode and bit-rate, the video frame rate, the resolution size, and the resultant stream bit-rate may be adjusted.

The application that encodes the multimedia content will have a corresponding application on the client system that will decode that received multimedia content. The different client applications available for viewing multimedia applications on a

multicasting address are Quick Time, Windows Media Player, and Real Player. These applications are free for download from the respective web sites. There are some proprietary applications that are available. Some of these applications are Quick Time Pro, Mpegable Player, and client Windows Media Player plug-ins that come with the server application. The proprietary client applications require a fee or license agreement. The application on the client side will not be discussed for the thesis. The type of coding and decoding (CODEC) that the application uses will vary. The main CODEC standards used for streaming multimedia are MPEG1, MPEG2, MPEG4, AVI, MOV, WMP, and MP3.

D.   MULTICASTING

There are three different addressing types for network traffic – unicast, broadcast, and multicast. Unicast addressing sends a data packet to another, single computer across a network. Unicast addressing is the dominate addressing method of the Internet Protocol (IP). Each connection between computers creates an individual data stream. If there are two computers that want to communicate to one computer or vice versa, there must be two separate data streams created. While unicast addressing is a one-to-one addressing scheme, broadcast addressing is a one-to-all scheme. Broadcast network traffic is sent to all computers that are connected to the network, whether every computer wants or needs to receive the data. This is an effective means of sending update messages, information messages, or other critical messages to all clients. There are security implications, such as a hacker denying service to the client computers on the network by flooding the network with broadcast messages and not giving the other computers on the network a chance to send out their packets. Such security issues are not within the scope of this thesis, however. Broadcast network traffic uses one data stream to all users. [WITT99]

Multicast network traffic is sent to a group of computers, but not all computers, on the network. Only the computers that want the data receive it. As a one-to-many, many-to-one, or many-to-many addressing scheme, multicast addressing combines the best of both unicast and broadcast methods. Multicast data may cross over to multiple networks, if necessary, and given the network devices support multicasting. Multicast uses one data stream from the sender to all the computers that want the data stream.
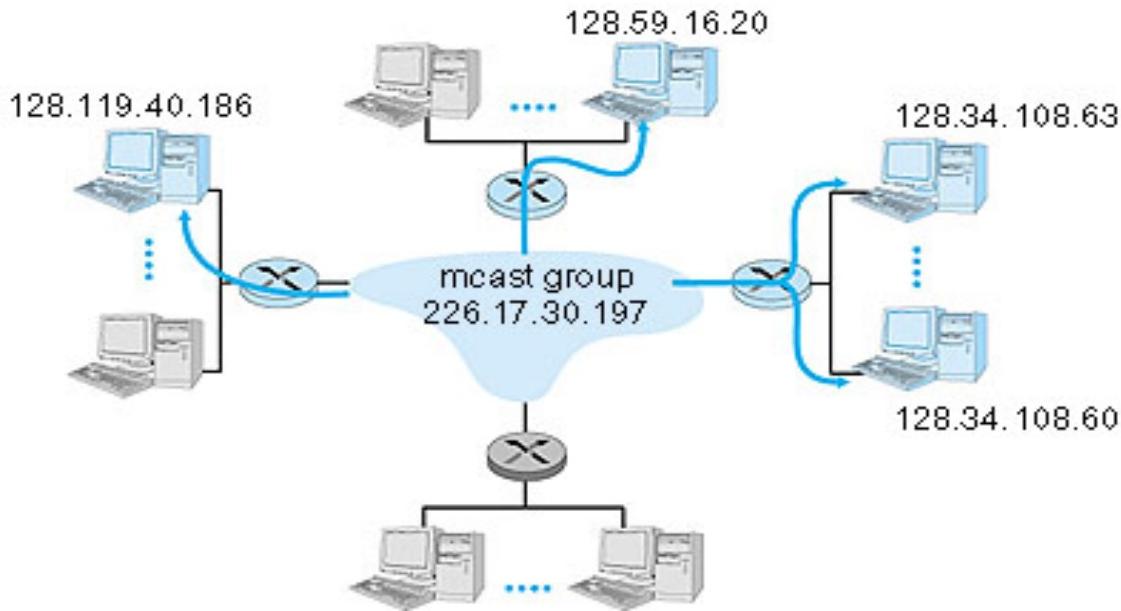
Figure 6.   Multicast Group [KURO01]

1.   Addressing

The Internet Engineering Task Force (IETF) dedicated the Class D address spacing for the use of multicasting.   The multicasting addressing space ranges from 224.0.0.0 to 239.255.255.255.   All 32 bits are used to specify a particular group; however, the first four bits are immutable, leaving only 28 bits to differentiate the groups. Since multicasting is a group address, there is not a need for a host address. [MILL99]

	Class A	*Class B*	*Class C*	*Class D*	*Class E*
Network Address Bits	8	16	24	32	32
Host Address Bits	24	16	8	N/A	N/A
Address Range	0.0.0.0 – 127.255.255.255	128.0.0.0 – 191.255.255.255	192.0.0.0 – 223.255.255.255	224.0.0.0 – 239.255.255.255	240.0.0.0 – 247.255.255.255
Subnet Mask	255.0.0.0	255.255.0.0	255.255.255.0	N/A	N/A
Addresses Begin with:	0	10	110	1110	1111
First Byte Values	0 – 127	128 – 191	192 – 223	224 – 239	240 – 255
Number of Networks	127	16,384	2,097,151	N/A	N/A
Number of Hosts	16,277,214	65,534	254	N/A	N/A

Table 1.   IPv4 Address Classes [ZACK01]

19

There are multicasting addresses that are used for special purposes or applications (see Table 2). The final entry in Table 2 is the multicast unregistered address space or local scope address range. This multicast address range is used within a local or private network and routers do not forward packets addressed to this range outside the originating network.   The IANA (http://www.iana.org) maintains a complete up-to-date list of all reserved addresses. [MILL99]

Applications	IP Address
All systems on the subnet	224.0.0.1
All routers on the subnet	224.0.0.2
All DVMRP routers	224.0.0.4
All OSPFv2 routers	224.0.0.5
All OSPFv2 routers	224.0.0.6
All RIP2 routers	224.0.0.9
All PIM routers	224.0.0.13
IGMP	224.0.0.22
Router-to-Switch	224.0.0.25
Network Time Protocol	224.0.1.1
Microsoft and MSNBC	224.0.12.0 - 224.0.12.63
Walt Disney Company	224.0.19.0 - 224.0.19.63
Local Scope	239.255.0.0 - 239.255.255.255

Table 2.   Assigned IPv4 Multicast Addresses

The latest version of IP, Version 6 (IPv6), is still in the deployment phase.   It will not be discussed in the thesis. See the IETF's Request for Comments (RFC) 2460 for more detailed information on IPv6.

2.   Protocols

The networking protocols are a set of rules used to allow the different networking components to communicate and route traffic. The protocols are defined at each layer of

the OSI model to allow communication between the layers. Multicasting requires new protocols, and new routing algorithms. These protocols must define ways to join or leave a multicast group.

Multicast protocols work at either the Data Link Layer (OSI layer 2) or the Network Layer (OSI layer 3). Data Link Layer multicasting can be used to transport the data over an individual link; however, there are problems due to the point-to-point connection orientation of the link layer and the incapacity to handle the one-to-many connection. These problems are mitigated by multicasting at layer 3, using logical addresses.

Network-layer multicast, or IP Multicast, is broken into two classes, Any Source Multicast (ASM) and Source Specific Multicast (SSM). ASM was the initial type of multicast developed and is still the primary form of multicast in use today. The basic functionality of Any Source Multicast is that the multicast IP address identifies the multicast group. Clients send the data to the multicast group address as the destination address. This allows any client to send data to any multicast group address on a network. Source Specific Multicast provides a group address for the multicast data stream and a unicast source address for the multimedia stream originator or source. The basic methodology of SSM is that all clients in the session can receive the data stream and only a few pre-specified computers (usually a server) in a session are allowed to add content to the session. [MILL99]

With unicast routing, there are two routing techniques used, with different variations of each. The two techniques are distance vector and link state. Distance Vector protocols generally choose a dedicated path that may or may not be the shortest distance. An early version of the distance vector technique is the Routing Information Protocol (RIP). Link State protocols became more popular because protocols such as the Open Shortest Path First (OSPF) allowed for quicker convergence over dynamic networks.

The multicasting routing techniques are derived from the two basic unicast routing protocols. The multicast routing protocols allow the network to find various multicast routes. The three main routing protocols used for multicasting are dense mode,

and sparse mode, and shared tree. Dense Mode multicast routing is used when network bandwidth is not constrained and the group members are in a concentrated area. The dense mode protocol uses either the Distance Vector Multicast Routing Protocol (DVMRP) or the Protocol Independent Multicast – Dense Mode (PIM-DM) which are based on the unicast distance vector routing protocol. Sparse Mode multicast routing is based on the unicast link state, best suited for widely dispersed members on a wide area network, and uses Multicast Open Shortest Path First (MOSPF). The shared tree is derived from the sparse mode and uses Protocol Independent Multicast – Sparse Mode (PIM-SM) and Core-Based Tree (CBT). [WITT99]

The routing protocols determine the routes the data follow from one location to another, but there needs to be a way to join, manage, and leave the multicast groups. This is accomplished through the Internet Group Management Protocol, which is up to version 3 (IGMPv3). There are different features in IGMP that allow for multicast group management. Part of IGMP is to send messages from the routers to the hosts and vice versa. The host will send a JOIN message to the router which will graft the traffic to the required route. If the host sends a LEAVE message, the router will prune back the route being used if no other user on the network desires the traffic. The router will also periodically check to see if there are any hosts still using the multicast group. If there are no hosts using the group, the router will prune back the route being used. For a more in-depth comparison and analysis of the various multicasting protocols, see the thesis *AN EVALUATION OF THE NETWORK EFFICIENCY REQUIRED IN ORDER TO SUPPORT MULTICAST AND DISTRIBUTED LEARNING NETWORK TRAFFIC* by Christopher Quick.

3. Multicast Advantage

Multicasting seeks to gain the benefits of both unicast and broadcast network traffic and while avoiding their weaknesses. It uses the targeted address scheme of unicastinig and the single stream benefits of broadcast addressing. Even though there is only one IP address, this address represents a group of N clients but is not subject to the required bandwidth and connections of N users as described earlier. Multicasting uses the group communication bandwidth of broadcast by using only one data stream; however, the benefit is that multicast is only processed by the computers that request the

data packets. Multicasting saves network resources by transmitting only one data stream to the N requesting users instead of having to replicate the same data stream N times.

Another advantage that multicasting has over unicasting is scalability. For example, if there are only 10 users requesting a 5 Mbps unicast multimedia stream, a 100 Mbps network segment would be able to handle the 50 Mbps of traffic. If 10 more users decide to join the network and request the same 5 Mbps multimedia stream, the 100 Mbps network segment is now above the maximum capacity, considering network overhead. No other users would be able to send traffic on the segment. With multicasting, there is no limit to the group size for the multicast group because there is no change in the amount of network capacity being used, so multicasting can handle the scalability dilemma without problems. The same 5 Mbps multimedia stream would service all group members, regardless of how many subscribed to the group. [WITT99]

Multicasting also benefits the server by reducing its tasking load. With multicasting, the multimedia server does not have multiple clients each vying for their own multimedia stream from the one multimedia file. The multicast server will only have one stream to service, and that one stream will be sent to the multicast group. This saves resources on the server by placing a lighter process load on it and allows for multiple multicast sessions to be established and thereby servicing multiple multicast groups.

MULTICAST TRANSMISSION

1 File and 1 copy

Transmission Source

1 copy

Switch

1 copy

Router

This diagram simulates a multicast network transmitting 1 file to 6 users. The original file is replicated once along the entire network path.

1 copy

4 copies

Router

Switch

Wide Area Network

1 copy

2 copies

Router

Switch
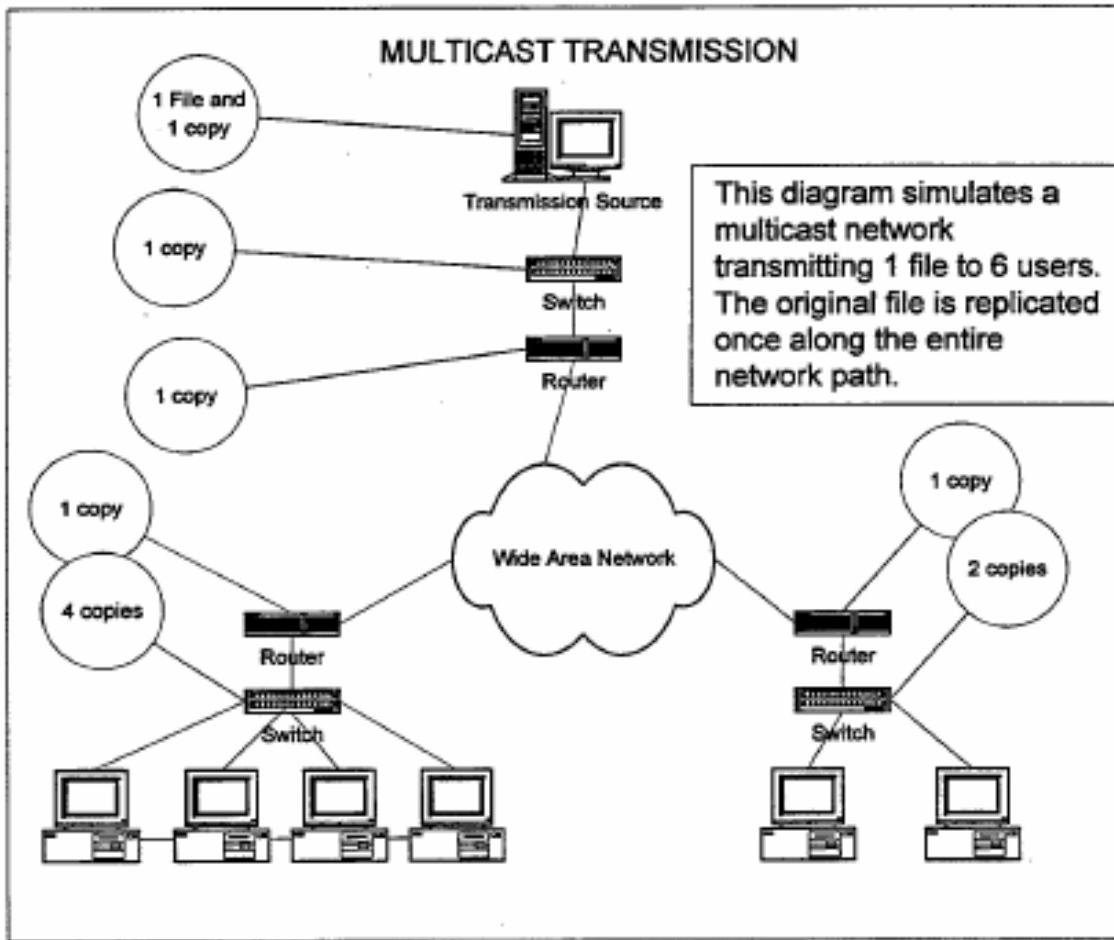
Figure 7.   Multicast [FLOU99]

4.   Multicast Applications

A multicast application is an application that has the capability to transmit and/or receive data from an IP multicast address. Multicasting has many benefits when it is used to its full potential, yet, there is a low demand for multicasting applications because there is still sufficient network capacity to handle the existing unicast network traffic of corporate enterprises.

There are several different applications that take advantage of multicasting. These applications are categorized into four basic categories: real-time multimedia, non-real-time multimedia, real-time data, and non-real-time data. Within each group, there are subcategories or specialized programs that handle the specific areas. Streaming multimedia, virtual reality, news and stock feeds, and data delivery are some of the programs. Although there are four basic categories of multicasting, the applications fall

under one of the two multicasting classes of either Any Source Multicasting (ASM) or Source Specific Multicasting (SSM). [MILL99]

Real-time multimedia applications are best used in today's Internet infrastructure and represent the greatest area of concentration. The category consists of, but is not limited to, video servers (SSM); live multimedia events (SSM); video conferencing (ASM); streaming audio (SSM); and distributed learning (ASM). The multimedia streaming server application that uses multicasting is the focus of this thesis. There are a variety of venders that provide multimedia streaming servers. Non-real-time multimedia applications are best suited for providing content that is recorded and then served to a multicast group (SSM). Most of the equipment used in real-time multimedia is also used for non-real-time multimedia and is part of the research in the thesis. [MILL99]

Real-time data applications are used to provide live news feeds (SSM), financial data from the stock exchange (SSM), or virtual-reality interactive gaming (ASM). Non-real-time data applications are used in providing data content delivery for servers, clients, databases, software distributions such as program revisions and updates (SSM), or synchronized shared databases (ASM). [MILL99]

5. Implementation Problems

Part of the problem with implementing multicasting is to get the personnel that manage and maintain the network to accept a new technology that potentially has many benefits but is unproven in a real-world environment. Many corporations, businesses, and industry are unwilling to take one step further from a lab environment and implement the technology into their business network. The new ideas are left to the colleges and universities to explore. There is a reluctance to accept multicasting because corportations are operating at network capacity and adding new multimedia content will increase network traffic beyond that capacity. Most technicians and network planners do not understand multicasting and there is a resistance to change what works; using the popular phrase of "If it's not broken, don't fix it." [MILL99]

Early implementations of multicasting caused problems in the installed network base. Most of the routing and switching network hardware was unable to handle the multicast traffic appropriately and the multicast traffic was either dropped or forwarded

as a broadcast message to all clients. There was also the problem of managing delays in networking where real-time data is a critical importance. The underlying routing did not offer a Quality of Service (QoS) management capability by guaranteeing bandwidth to the server. These problems hindered the acceptance of multicasting. [MILL99]

Early many-to-one multicasting had an implosion problem. It happens when all the group members send an acknowledgement back to the source of the multicast traffic. These acknowledgements often flooded or overwhelmed the server. As a result, the many-to-one multicasting became phased out and Source Specific Multicasting and Any Source Multicasting became more prevalent. [WITT99]

Although the application is multicast capable, there are similarities to unicast applications. The session for the multicast application must be set up and then torn down; the multicast groups must be established, publicized, and managed; and multicast addresses must be managed by a network technician to ensure that multicast group addresses do not overlap. There is no standardization among the different applications between the vendors, and the multicast protocols are different as well. The client application that is used by one vender may or may not work with the multicast stream from another vender, which raises issues such as cost impacts and compatibility of operating systems. [WITT99]

Security can be an issue for multicasting, which has also caused problems with acceptance. The multicast group address may be publicized in advance to provide the group members with advanced notice to receive the data. The information may be available to the public and anyone who may want to receive the data can access the data stream. Some of the information may not be releasable to the general public because of the proprietary rights on the data. If no encryption is affected on the data or multimedia stream, this becomes a concern in so much as the proprietary information may be compromised. Additionally, a hacker could transmit on a multicast address because there are no restrictions on the users in the group or the timing for which a client has permission to transmit. Even if the data being sent is encrypted, the hacker may still perform a denial of service (DoS) attack against the multicast group address. For a more

wide spread acceptance of multicasting, security concerns must be more aggressively addressed. [MILL99]

The basic background necessary to understand the thesis has been discussed. The terminology covered gave a brief description of the various pertinent areas of networking, multimedia, system components, and multicasting. The streaming multimedia terminology and formats, system composition, and multicasting benefits should be clear. The next chapter discusses the basis of the metrics used to analyze the multicast server.

THIS PAGE INTENTIONALLY LEFT BLANK

III. METRICS USED IN COMPARISON

A. BASIS OF METRICS

1. Introduction

This chapter will discuss the details of how the metrics will be determined for comparing the various platforms that are being tested. The thesis is on the methodology of using the metrics for the comparison of various hardware and software platforms that do not have a lot of similarities, not for the minimum requirements to achieve the metrics. The equipment used is a sample case study that aided in defining the methodology of the metrics to be used for comparison. The metrics that are going to be covered in this chapter come from eliminating areas that are different and could not be measured side by side. The metrics that were finally settled upon for testing are areas where there is a commonality between the various hardware and software component configurations, such as networking or video quality characteristics. There are several common areas that are then broken down into subcategories that have measurable metrics. These metrics are then combined to give a rating which will aid in determining which system has the best performance to price ratio.

There are many tests available to determine which piece of hardware is better than another piece of hardware. When comparing different types of microprocessor units such as the Intel Pentium IV to the AMD Athlon XP, the basic metric that can be used is how many instructions it can execute in a given amount of time, whether it be simple arithmetic operations, swapping of memory locations, or a basic logical comparison. A more in-depth discussion can be found in *Computer Architecture* by Hennessey and Paterson. When comparing video cards, there are several basic metrics that can be used to compare the cards. These metrics include, but not limited to: 1) how many pixels can be drawn in a given time; 2) number of colors; 3) brightness; 4) contrast; 5) frames per second; and 6) 3D capabilities. The comparison of hard drives is based on storage size, the interface, access time (a composite of seek time, rotation speed of the disk platter in revolutions per minute (rpm), and data transfer speed), and buffer memory. Network cards are compared on the connection speed, buffer memory, and transfer rates.

These individual components have distinct metrics that can be used for comparison to each other. When all the components are integrated, the system might be compared based on the individual components and individual metrics, however, metrics that test the complete systems may be more appropriate and illuminating. These tests may include how fast a particular application will open a document or complete a series of mathematical computations or displaying of graphical tests. This will produce a benchmark or baseline that can be used in comparison with other similar systems. [CASE03]

2. Comparing Different Hardware Platforms

The comparison between similar components is fairly straight forward and easy to accomplish. When comparing a microprocessor that processes Complex Instruction Sets or Reduced Instruction Sets, it is easier to determine which microprocessor is faster based on the number of instructions that were processed in a given amount of time. When comparing video processing cards, it is easy to compare the metric of frames per second because that is something that is understood by all manufacturers. The difficulty is trying to compare hardware components where the items are different, using different metrics, but basically accomplish the same task. The comparison of a Complex Instruction Set Microprocessor (CISC) to a Reduced Instruction Set Microprocessor (RISC) is not easily accomplished. With a CISC processor, the instructions must be converted, reduced, or broken into smaller instructions before that entire instruction is processed. Once this CISC instruction is reduced to its lowest level, can it then be compared to a RISC processor that is basically the equivalent group of instructions? Once the basic group of instructions has been determined, then the comparison can be made of how many instructions are processed in a given amount of time. But even then it may not be equal because of the way the microprocessor handles a given instruction. The design of a system can determine different variables as well. Certain systems may use high speed memory such as PC2700 while others may use memory that is of average speed like PC133. The amount of cache used can make a difference in the speed of execution due to the decrease in loading time or pipeline staging of instructions. The interface speed of the hard drive may be faster in one system over another or different altogether. Various interfaces such as the Small Computer Systems Interface (SCSI), Enhanced Integrated

Drive Electronics (EIDE), and Serial Advanced Technology Attachment (SATA) all have different interfaces and different speeds in which data is transferred to memory. The way graphics are drawn to the screen makes a difference in the visual representation of what is happening inside the computer, especially when it comes to graphics that draw in three dimensions. [CASE03]

To compare systems that are of a different design, operating system, and manufacture requires some common denominator to differentiate the performance characteristics. These common factors are what will be identified for the methodology of testing and comparison of the multimedia servers.

3. Comparing Different Operating Systems

Not only does the hardware have many variables in the way metrics are determined for comparing components, but operating systems are designed differently as well. There are several different operating systems available, and each operating system has points that are stronger than others, partially due to the way the operating system handles code and performs various tasks. The major operating system venders are Apple, Mandrake, Microsoft, RedHat, and Sun. Each of these operating systems handles tasks differently. The different areas that vary in an operating system include, but are not limited to: 1) the way it manages memory, 2) how it schedules tasks to be executed by the processor, 3) the accessing of hardware, whether directly or indirectly; 4) transferring data from main memory to secondary storage, 5) handling of code for graphics drawing, and 6) the communications protocol stack used to access the physical network. All of these areas play a role in the overall efficiency of the operating system. This thesis will not quantify which operating system is better in the different areas; there are many other reports, reviews, and tests available for those results. Rather, this thesis is looking at an area that is common to all the operating systems and that can be used for a comparison of the different servers that are available for multicasting multimedia content. The main area in the operating system that the thesis is looking at is the ease of use, programmability, and flexibility of the operating system to accommodate the necessary programs, functionalities, and hardware required to provide streaming multimedia content via multicasting. The minimum requirements for the operating system are specified by the vender for their system and will not be discussed in this thesis.

## B.	DETERMINING THE METRICS
### 1.	Using Common Metrics

To determine the metrics that are going to be used in the server comparison, there has to be a common denominator established, and one that can be measured in some form or fashion. The common areas apparent for comparison of the hardware and software servers are in the areas that deal with the network capabilities and capacity and the multimedia capabilities. One of the main factors in a network is the bandwidth capacity. How much data can the network handle at one time, how much capacity does the server supply? How many distinct, simultaneous connections or multimedia multicast streams can the server support? Common areas in multimedia include frames per second, resolution quality, and the sampling rate of the audio. Another major area that is looked at is the area of serviceability, such as the equipment upgradeability and expandability, as well as the flexibility of software programming and the services provided.

### 2.	Disregarding Areas That Are Not Directly Comparable

Since there are many areas in this study that can not be compared due to the difference in technology, those areas will not be covered or compared. There are many reports, companies, web sites, and test that are available for specific graphic cards, hard drives, CPUs, and other individual components. This thesis is not about testing individual components; it is about defining a methodology and looking for a common factor for comparison of different systems. The main purposed behind the comparison is to try and identify areas that are compatible for comparison. Just as it is difficult to compare vegetables and fruits, and within fruits, apples and oranges, so it is in comparing hardware and software with different areas of technology. However, there are some areas that the study has found to be in common among the different servers. All individual hardware component comparisons will be considered in the analysis, which includes items such as processor(s), memory, disk storage, and graphics adapters. Nor will general comparison of the operating systems be included. Any comparison of operating systems will be with regard to there suitability for streaming multicasting multimedia content, not which one is better.

C.   METRICS USED FOR STUDY

1.   Streaming Capacity

This metric will be evaluate the number of multicast streams the server can support without showing any degradation in the quality of the individual multimedia stream that is being delivered to the end user. Packet loss is a normal occurrence in networking, although it must be minimized as best possible. This metric measures the server's tolerance of packet loss and identifies when it starts to degrade in such a manner that the end user notices flicker or jumping in the video, breaks in the audio stream, or large losses of packets. This measurement will be a number based on actual streams being distributed.

2.   Bandwidth

This metric will be used to determine how large a multicasting stream is generated from the multimedia that is being distributed by the server. Part of the stream size will come from the overhead that is generated from the application program that is distributing the multicast stream and part of the stream size will be determined by the compression algorithm that is used by the generating server application. Overall, what is being looked at is the total size of the multicast multimedia stream being generated. This measurement will be in Megabits per second (Mbps) for the same quality of audio and video (i.e. frame rate, resolution, and audio).

3.   Quality of Audio/Video

This metric will be used to determine the quality of the audio and video that is being streamed out to the multicast address. Part of the metric will be looking for any degradation of the audio or video. This measurement will be broken down into subcategories based on the following areas: frames per second, resolution size, coding format, and audio sampling rate.

a.   Frames per Second

This metric will be used to determine how many frames per second are used to record, generate, or display the video. The standard high quality video range is between 24 and 30 frames per second. The next level of video that still has some discernable quality requires between 15 and 24 frames per second. This range is good for

video where there is little movement in the captured video. In any recording that provides less than 10 frames per second, the video is of poor quality and noticeably jumpy. The two thresholds used will be 30 fps or 15 fps.

b. Resolution Size

This metric will be used to compare the size of the resolution for the video that is being recorded and/or played back. There are different formats in which the media may be captured. There is the NTSC resolution for standard television and the HDTV resolution for digital television / digital video.   The sizes will be 720x480 (DV 16:9 Aspect Ratio), 640x480 (NTSC 4:3 Aspect Ratio), 320x240 (½ NTSC 4:3 Aspect Ratio), and 160x120 (¼ NTSC 4:3 Aspect Ratio).

c. CODEC (Coding / Decoding) Format

This metric will be used to assess the format by which the multimedia stream is encoded. There are several major formats that are being commonly employed: 1) QuickTime, .avi or .mov; 2) MPEG1; 3) MPEG2; 4) MPEG4; 5) H.261 or H.263; and 6) RealAudio / RealVideo. The baseline for this category will be whether or not the server supports MPEG1, MPEG2, or MPEG4.

d. Audio Sampling Rate

This metric will be used to contrast the sampling rate of the audio being generated for the multimedia stream.   There are several different levels of audio recording and sampling rates that are being used for broadcasting in a digital format. The audio could be in stereophonic, using two audio tracks, or monophonic, using just one audio track. The sampling rate for audio has a wide range going from a monophonic quality of 8 Kilobits per second to better than CD Stereo quality of 192 Kilobits per second. The higher the sampling rate, the better the quality of the audio, but the size of the file and bandwidth required for streaming will increase proportionately, as well. The measurement used will reflect whether the output is stereophonic or monophonic quality.

4. Maintainability

This metric is broken down into several categories that define areas of actual cost of the hardware, ease of maintenance and availability of support from the manufacturer,

and the ability to upgrade the hardware with additional components, such as more memory, disk space, network interfaces, and processors.

a.   *Scalability*

This metric will be used to indicate if the server has expansion capabilities for adding more individual hardware components, either memory, hard disk storage, or network cards. Can the equipment be upgraded to handle more growth requirements? Does the network connection speed support 10 Mbps, 100 Mbps, or 1000 Mbps? Can it be equipped to support faster transfer speeds? This value of this measurement will be either "Yes" or "No."

b.   *Cost*

This metric will be used to indicate the price of the equipment that is being tested. This measurement will use the following rating: Expensive, Affordable, and Budget. To receive an Expensive rating, the cost of the equipment must be over $5,000. To receive an Affordable rating, the cost of the equipment must be between $5,000 and $1,500. To receive a Budget rating, the cost of the equipment must be less than $1,500.

c.   *Maintenance*

This metric indicates if the equipment is supported, at additional costs to the end user, through a service contract or if the equipment is supported only for hardware failures but the user must maintain the system and perform upkeep. It will reflect the required maintenance to keep the equipment operational.   The following questions are addressed by this metric: Does the equipment require the equipment to be sent back to the manufacturer for corrective maintenance or can a local technician repair the equipment? Can the equipment be upgraded through Commercial-Off-The-Shelf (COTS) equipment or is it proprietary and parts can only be purchased from the original manufacturer? Is the software maintained by the end user or does it require intervention by the product developer? The value of the measurement will be either "COTS" or "Proprietary."

5.   Versatility

This metric is used to measure the programmability of the hardware and software. Does the equipment do more than just multicasting multimedia content or are there other

applications or services available that can be run simultaneously? What effect do the other services have on the multicast multimedia streaming? Will the equipment require constant administration by a user or will the equipment function autonomously? Does the equipment require periodic updates of software and hardware or will the equipment function as originally delivered without the necessity of updates. While there is expected to be a minimum of some form of human intervention, the goal is that it should require minimal amount of effort, research, knowledge, hands-on time, and maintenance to keep the server operational and functional. The multimedia multicast streaming server is used as a distance learning tool and will usually be operated by one trained in server administration or a certified technician. The measured value will be based on the following rating: High Versatility / High Maintenance (H/H); High Versatility / Low Maintenance (H/L); Low Versatility / High Maintenance (L/H); or Low Versatility / Low Maintenance (L/L). To receive a measurement of High Versatility / High Maintenance, the equipment requires an administrator trained in maintaining the equipment on a daily basis, but the equipment can be used for additional services. The High Versatility / Low Maintenance rating will be given to equipment that can be used for multiple services, but requires little intervention once the services are established. The Low Versatility / High Maintenance will be used for equipment that can only provide multicasting multimedia streaming but still requires daily technical intervention to operate. The Low Versatility / Low Maintenance will be for equipment that only provides multicasting multimedia streaming and minimal operator intervention to provide the streaming content.

6.   Security/Encryption

This metric is used to determine if there is any security or encryption capability for the transmission of the multimedia content. Some media may be copyright protected or it may be inappropriate for the content to be freely distributed. Without some form of protection, the multimedia content could easily be captured and recorded by a variety of available programs. End-to-end encryption, provided by the application, will provide a layer of security against the casual eavesdropper or snooper. The assessed value for this metric will be "Yes" or "No."

D.   METRICS MATRIX

The defined metrics will be formatted such that they can be used easily to compare the different hardware and software servers. The metrics will be placed in a tabular chart to aid in checking-off the appropriate areas or recording the appropriate numbers. See Figure 8 for an example.

Metric \ Server		Server 1	Server 2
Streaming Capacity (#)			
Bandwidth (Mbps)			
Quality of Audio and Video	- Frames per Second	30 / 15	30 / 15
	- Resolution Size	F / ½ / ¼	F / ½ / ¼
	- Resolution Size	M1 / M2 / M4	M1 / M2 / M4
	- Audio Sampling Rate	Stereo / Mono	Stereo / Mono
Maintainability	- Scalability	Yes / No	Yes / No
	- Cost	E / A / B	E / A / B
	- Maintenance	COTS / Proprietary	COTS / Proprietary
Versatility		HH / HL / LH / LL	HH / HL / LH / LL
Security		Yes / No	Yes / No

Figure 8.   Metrics Matrix.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. SYSTEM CONFIGURATION

A. SYSTEM FOR STREAMING

This chapter is the detailed description of the different multimedia multicast systems evaluated in the thesis. There are two different configurations that are used, recorded multimedia content and live multimedia content. The first configuration, recorded multimedia content, is used for providing a multicast multimedia stream that was previously recorded using a dedicated system to capture, edit, and convert multimedia content to a digital MPEG variant for streaming. This removes a large burden off the server processor. This process makes the multimedia multicasting server a multimedia file server; however, the system is sending the multimedia streamed format to a multicast address. This system is a Source Specific Multicast multimedia server. The second configuration, live multimedia content, is used for streaming live content that is converted to a digital MPEG variant in real-time to the multicast address. The process includes capturing the analog multimedia content, converting to the appropriate digital format, and streaming the multimedia content to a multicast address. The system places a heavy load on the processor for conversion and encoding; and is a Source Specific Multicast multimedia server. Following the description of each multimedia multicast system evaluated is why each system was used in the testing.

1. Streaming Live Content

The multimedia multicast server is capable of streaming live multimedia content to a multicast address. The input is a camera that is capable of delivering a composite video signal and stereo audio channels (RCA) to the computer for conversion to a digital data packet for streaming. The camera details will not be covered; there are hundreds of cameras to choose from that have S-Video, Digital Video, and/or Composite Video as output. Although the content is not live, lectures or training tapes fall under this category because the content is being played from an external device such as a Video Cassette Recorder/Player (VCR) or Digital Versatile Disc (DVD) player and converted in real-time. For more information about recording multimedia content in a classroom environment, see *Streaming Media Bible* by Steve Mack.

When streaming live multimedia content, there is a delay from the component input (camera, VCR, DVD player) through the server to the network due to the digitization encoding, and the encryption process, if any is performed. The delay is minimal and will be compounded as the multimedia stream traverses the network through the routers and switches. This delay can be minimized on the network with the routers and switches configured properly and providing Quality of Service (QoS). The client computer then decodes the multimedia content.

2.   Streaming Recorded Content

Some multimedia multicast servers have storage capabilities for storing multimedia content that has been digitized using a variety of different programs. Making the digitized content available for streaming to a multicasting address, it must go through a conversion and optimization process for streaming; adding hinted tracks that coordinates the timing of audio and video. There is a plethora of information on how to convert analog multimedia content to digitized multimedia content. For more information, see *Streaming Media Bible* by Steve Mack.

B.   EVALUATED SYSTEMS

1.   Apple XServer

This is the detailed description of the Apple XServer used during the testing of the multicasting multimedia server. The Apple XServer was used in a live configuration format. The same configuration was used in a streaming recorded format. The streaming content was previously recorded through alternate means.

The hardware configuration as tested:

> Dual G3 Processors @ 1.3 GHz
> 512 MB DDR333 RAM
> 60 GB ATA133 Hard Drive
> CD-ROM Drive
> Dual 100/1000 Mbps Network Interface Connection
> ATI VGA Graphics Card
> FireWire 400
> FireWire 800
> USB 1.1 Connection

The operating system configuration as tested:

> MAC OS X Server (Version 10.2.6)

Additional services such as ftp, mail, and web disabled.

Software used for live streaming:

QuickTime Broadcaster

Software used for recorded streaming:

QuickTime Streaming Server

Price of system as configured:

$3,899 for Server Hardware, MAC OS X, and Streaming Applications
$999 for Server support services
$950 for 3 years onsite hardware service, telephone and e-mail support

Apple Computers was asked to participate in the multimedia multicasting server comparison for the thesis. They offered to provide the Apple XServer for the thesis study of the comparison of multicasting multimedia servers. This is the configuration they were able to provide for testing.

2.   Video Furnace Server

This is a detailed description of the Video Furnace Server that was used during the testing. The Video Furnace Server was only tested in a live streaming environment. The future released versions will incorporate a recorded streaming capability.

The hardware configuration as tested:

Pentium 4 Processor @ 1.8 GHz
512 MB DDR266 RAM
82 GB EIDE Hard Drive
CD/DVD Drive
100/1000 Mbps Network Interface Connection
VGA 32 MB Graphics Card
PCI Audio/Video Capture Card
IEEE 1394 Firewire Connection
USB 2.0 Connection

The operating system configuration as tested:

RedHat Linux (Version 7.3)
Additional services such as ftp, e-mail, and web disabled.

Software used for live streaming:

VF LiveFurnace – real-time encoding server software
VF InStream – client viewer application
VFAgent – supports the automatic distribution of content viewer to clients
VF OnGuide – electronic programming guide
VF Pilot – GUI interface for configuring LiveFurnace

Price of system as configured:

$12,500 for Server Hardware and Live Encoding/Streaming Applications
$2,000 for upgrade support services
$2,500 for License Server (One time fee)
$8 - $12 per year per concurrent viewer license

Video Furnace was asked to participate in the multimedia multicasting server comparison. Video Furnace Inc. offered to provide a server for the thesis study with the possibility of a future purchase by the Naval Postgraduate School IT Department. This is the product configuration provided.

3.   VBrick 3200

This is a detailed description of the VBrick 3200 that was used during the testing. The VBrick 3200 was only tested in a live streaming environment. The multimedia content is inputted into the unit from a VCR, DVD Player, or any video camera that has a S-Video or composite output. There is no storage capability in the VBrick 3200; however, VBrick does offer other solutions with storage capabilities.

The hardware configuration as tested:

Standalone Unit
S-Video or Composite Audio/Video Input
Audio/Video Encoder
10 Mbps Network Interface Connection

The software configuration as tested:

Firmware Version 2.5.1
VBAdmin – Remote Administration Utility Program

Price of system as configured:

$3995

This server was used in the comparison because the Naval Postgraduate School IT Department had the unit for a different project. The VBrick 3200 was configured for unicasting due to multicast problems. The problem was researched for the thesis *AN EVALUATION OF THE NETWORK EFFICIENCY REQUIRED IN ORDER TO SUPPORT MULTICAST AND DISTRIBUTED LEARNING NETWORK TRAFFIC* by Christopher Quick to correct multicasting deficiencies. VBrick was asked to participate in the multimedia multicasting server comparison. Several methods of contact were

attempted with no response. The author was unsuccessful in gaining additional equipment or support for the thesis comparison from VBrick, Inc.

4. General Purpose Server

This is a detailed description of the general purpose server used during the testing. The general purpose server was only tested in a recorded streaming environment. There are several applications that are available for recorded and live streaming; however, there was only one trial version available. The venders would not offer a working copy of the software application available for the thesis testing. This server was used for testing.

The hardware configuration as tested:

Dual Pentium 3 Processors @ 500 MHz
768 MB PC133 RAM
80 GB SCSI Hard Drive
CD/DVD Drive
100/10 Mbps Network Interface Connection
32 MB VGA Graphics Card
USB 1.1 Connection

The software configuration as tested:

Windows2000 Server
Additional services such as ftp, e-mail, and web disabled.

Software available for live streaming:

MPEGable Broadcaster Trial Version by Dicas of Germany

Software used for recorded streaming:

StreamPump Free Version by VBrick, Inc.
Darwin Streaming Server by Apple, Inc.

Price of system as configured:

$899 for computer hardware
$899 for Windows 2000 Server

This server was used in the comparison because the Naval Postgraduate School had the unit on campus for a different project. The general purpose server is used as a generic Windows2000 server for the networking test lab. This server provided a test platform used as a competitive operating system to Apple and RedHat.

After many unsuccessful attempts to have various venders participate, this is the list of equipment available for the comparison testing.

THIS PAGE INTENTIONALLY LEFT BLANK

V. SYSTEM TEST RESULTS

A. COMPARISON RESULTS

This chapter discusses the results of the metrics performed on each system. Additional items are addressed pertinent to each server that does not fall under the categories described in the metrics matrix. The system metrics matrix follows the results description. The system metrics matrix has a column labeled "Score". The column displays the running total of positive and negative marks. Each factor is weighed equally. The last row/column has the final score of the evaluated server.

1. Apple XServer

The following is the metrics statistics for the Apple XServer. For the streaming capacity of the XServer, there is one live stream. The QuickTime Broadcaster application can receive one input and send one output. Broadcaster has the capability to record the live multimedia stream for playback at a later date and time. For recorded multimedia content, the number of multicast multimedia streams that were able to be transmitted was limited to the network bandwidth of the installed network interface card. Although the number of streams is limited by network bandwidth, that limit was approximately 60-65 multicast streams. The bandwidth required for streaming the multimedia content to the multicasting address will vary depending of the recording application program. The live multimedia stream rate can not be set; but at 30 fps, 320x240, 44 KHz stereo sampling at 128 Kbps, and MPEG4, the stream rate was 1.5 Mbps. The range varies from 800 Kbps to 3.5 Mbps. For the recorded content, QuickTime Streaming Server would only stream MPEG4 without additional plug-ins. The bandwidth required to stream the recorded MPEG4 content was based on the recorded rate of the content. The required bandwidth ranged from 500 Kbps to 2 Mbps. The quality of video allowed for capture up to 30 frames per second. The resolution size was restricted to 320x240 (NTSC ½ Screen) or 160x120 (NTSC ¼ Screen). The CODEC format was MPEG4 unless there was a plug-in added to the server application which allowed for other CODECs, which cost an additional fee. The audio sampling rate ranged from 8 KHz to 48 KHz and had the option of mono or stereo recording. The

camera used for capturing the live content needs FireWire output to work with the XServer. The Apple XServer provides scalability for adding more memory, hard drive storage, and network cards for greater network bandwidth, although there is sufficient network bandwidth for streaming multimedia content to a multicast address. The cost of the server is affordable; however, that could be changed after adding in support services, on-site maintenance, and the necessary plug-ins to stream older MPEG formats or application required to convert MPEG1 to MPEG4. Unfortunately, the maintenance is proprietary, although some of the upgradeable components may be purchased by Commercial-Off-The-Shelf (COTS) venders. The XServer offers high versatility allowing the server to provide multiple services such as ftp, e-mail, and web pages; but this will reduce the recorded streams provided due to reduced network capacity consumed by other services. Although the XServer offers high versatility, it also requires a high amount of administration to offer multimedia streaming content, as described below. The encryption security feature is not existent for security to the multicast multimedia streaming content. The XServer does provide all the other security features expected from an application server.

Here are some additional notes on the Apple XServer that need to be mentioned. There was some difficulty getting the multimedia streams looped internally through the local address relaying to a multicasting address. The XServer would buffer the multimedia stream before relaying to the multicast address. This may take a large amount of time depending on the size of the multimedia file buffered. For the clients to view the multimedia stream on the multicast address using QuickTime, the Session Description Protocol (SDP) file has to be captured from the directory where the Play List information is stored. The captured SDP is edited to change several parameters. The multicasting address and port numbers are entered to reflect where the multimedia stream is being distributed allowing QuickTime to receive the multicast multimedia stream. After that was accomplished, the file had to be placed in a public location on the intranet for the users to access, and finally the user would be able to view the stream. The same process must be accomplished for streaming live multimedia content. This requires excessive and unnecessary administration by the technician or administrator for the clients to view the multicast multimedia stream. For help not published in manuals,

Apple maintains and monitors a mailing list for the Apple XServer and QuickTime Streaming Server products where an individual can post questions for additional help which may be answered by someone who has had a similar problem in the past, or by an Apple Representative. The NPS MOVES Department has some MPEG1 and .mov multimedia clips from past lecturers that were used as sample for streaming via multicasting. The .mov files would not be streamed by QuickTime Streaming Server although they were hinted and optimized for streaming in the native format for QuickTime. The files had to be converted to MPEG4 embedding the hint tracks. Although Apple claims that QuickTime can do this, it must be the QuickTime Pro version, purchased for an additional fee.

Metric \ Server		Apple XServer	Score
Streaming Capacity (#)		Live – 1 Recorded – Limited by Network Bandwidth	+ / (1)
Bandwidth (Mbps)		800Kbps to 3.5 Mbps	+ / (2)
Quality of Audio and Video	- Frames per Second	30	+ / (3)
	- Resolution Size	½ NTSC	- / (2)
	- Resolution Size	MPEG4	+ / (3)
	- Audio Sampling Rate	Stereo	+ / (4)
Maintainability	- Scalability	Yes	+ / (5)
	- Cost	Affordable	+ / (6)
	- Maintenance	Proprietary	- / (5)
Versatility		High Versatility High Maintenance	- / (4)
Security		No	- / (3)
Final Score			3

Figure 9.   Apple XServer Metrics Matrix.

2.   Video Furnace Server

The following is the metrics statistics for the Video Furnace Server. The streaming capacity of the Video Furnace Server is one live stream. The audio/video capture-card and VF LiveFurnace application has one input and transmits one multicast stream. VF LiveFurnace does not offer the capability to record the live multimedia stream for playback and being incorporated in the next release. The bandwidth required for streaming the multimedia content to the multicasting address varies depending of the recording application. The stream selection rate range varies from 500 Kbps to 7 Mbps. The quality of video offers 30 frames per second. The resolution sizes include the options of: 720x480 (NTSC Full Screen), 480x480, 480x240, and 320x240 (NTSC ½ Screen). The CODEC format is MPEG1 with future release to provide a choice of MPEG1, MPEG2, or MPEG4. The audio mode offers: Mono, Stereo, Dual (2 mono), Joint Stereo 0 (3 KHz – 22 KHz), Joint Stereo 1 (6 KHz – 22 KHz), Joint Stereo 2 (8 KHz – 22 KHz), Joint Stereo 3 (11 KHz – 22 KHz) The audio bit-rate recording ranges from 32 Kbps to 384 Kbps. If the audio bit rate is less than 96 Kbps, it is recommend the audio input is set to monophonic instead of stereophonic. The camera used for capturing the live content must have a composite audio/video (RCA) or S-Video output to work with the VF LiveFurnace server that was tested, but Video Furnace does offer IEEE 1394 input solutions. Video Furnace provides scalability for adding more memory, hard drive storage, and network cards to increase network bandwidth. There is sufficient network bandwidth for streaming multimedia content to a multicast address using the 100 Mbps connection. The cost of the server is expensive; however, Video Furnace offers a high quality application that requires minimal administrator intervention. Just hook up the audio/video input, select the size of the multimedia stream, video resolution, audio sampling rate, and the server is off and running. The maintenance can be proprietary through Video Furnace (hardware/software upgrades) or maintained by the IT Department staff (hardware). The upgradeable components may be purchase by Commercial-Off-The-Shelf venders; however, the server is delivered with sufficient hardware capabilities. The Video Furnace Server offers high versatility allowing the server to provide multiple services such as ftp, e-mail, and web pages; but this is not recommended due to placing an added burden on the processor and reducing the network

capacity consumed by the other services. Adding to the high versatility of the Video Furnace Server, the server adds minimal administration required to operate and maintain the multimedia streaming content. A benefit offered by the Video Furnace server is security feature providing encryption to the multimedia streaming content. The encryption uses a 64-bit block cipher and a 128-bit key. The Video Furnace Server provides all the other security features expected from a Linux application server.

Here are some additional notes on the Video Furnace Server that need to be mentioned. Administration requires minimum effort by the technician or administrator to setup and transmit the multicast multimedia stream. Once the multimedia multicast stream is established and running, there is no additional requirement needed by the administrator, unlike the Apple XServer. The server offers the ability to edit the VF OnGuide to provide a description of each channel. If CNN is on one channel and a training video was on a second server/channel, the VF OnGuide reflects the broadcast of each channel. This would be beneficial when the Video Furnace is used for distributed learning and multiple classes are provided on the NPS campus network. Another application the Video Furnace offers is VF Pilot. The application allows remote monitoring and adjustments for each server.

Client access to the multimedia multicast stream requires the user open a web browser and entering the Universal Resource Locator (URL) to the server. The VF OnGuide application is transferred to the client allowing the user to select the channel. The selected channel is opened in the view window and the application will start receiving the multicast multimedia stream. The VF OnGuide application is simple to operate and resembles an OnScreen Channel Guide offered by cable providers. The operating system is not restricted; the VF OnGuide works with the three major venders. The VF OnGuide does not leave a footprint or data on the client computer when the application is closed.

Metric \ Server		Video Furnace Server	Score
Streaming Capacity (#)		1	- / (-1)
Bandwidth (Mbps)		Ranges from 500K to 7 Mbps	+ / (0)
Quality of Audio and Video	- Frames per Second	30	+ / (1)
	- Resolution Size	Full NTSC	+ / (2)
	- Resolution Size	MPEG1	+ / (3)
	- Audio Sampling Rate	Stereo	+ / (4)
Maintainability	- Scalability	Yes	+ / (5)
	- Cost	Expensive	- / (4)
	- Maintenance	Proprietary	- / (3)
Versatility		High Versatility Low Maintenance	+ / (4)
Security		Yes	+ / (5)
Final Score			5

Figure 10.   Video Furnace Metrics Matrix.

3.   VBrick 3200

The following is the metrics statistics for the VBrick 3200. The streaming capacity is one live stream from the VBrick. The audio/video capture component of the VBrick 3200 receives one input and transmits one output. VBrick 3200 does not offer the capability to record live multimedia stream for playback. The bandwidth required for streaming the multimedia content to the multicasting address varies from 1 Mbps to 3.5 Mbps. The quality of video offers the range from 2 to 30 frames per second. The resolution size offered is 352x240 (NTSC ½ Screen). The CODEC format output is MPEG1. The audio sampling rate ranges from 8 KHz to 48 KHz and a bit rate from 32 Kbps to 192 Kbps for stereo recording. The camera used for capturing the live content must have a composite audio/video (RCA) or S-Video output. The VBrick 3200 does not provide scalability for adding more memory, hard drive storage, or increased network

bandwidth, although there is plenty of network bandwidth available for streaming multimedia content via the 10 Mbps connection. The cost of the server is affordable; however, the price is high for providing one streamlined service. The Apple XServer and General Purpose Server offer versatility in the same price range. VBrick does not offer a robust server that can provide other services. The VBrick 3200 is optimized for encoding multimedia multicasting stream for the proprietary client viewer. To transmit the audio/video input, select the size of the multimedia stream, video resolution, audio sampling rate; and the server is off and running. The maintenance is proprietary through VBrick. There are no upgradeable or serviceable components in the VBrick 3200. The VBrick server offers low versatility providing one optimized service. Adding to the low versatility of the VBrick 3200, the server also adds the benefit of low administration required to operate and maintain the multimedia streaming content. The administration of the VBrick 3200 is accomplished through a computer serial interface using a Windows terminal application or by a unicast connection over the network. There are no security features that provide encryption to the multimedia streaming content. The only security feature offered by the VBrick 3200 is the administrator login requires a password.

Here are some additional notes on the VBrick 3200 server that need to be mentioned. Administration by the technician or administrator to transmit the multicast multimedia stream requires minimum effort. Once the multimedia multicast stream is established and running, there is no additional administrator requirement. The default setting for streaming is set for unicast. The VBrick Admin application must be installed on a remote Windows PC which allows configuration of the VBrick 3200. The configuration is set through the serial interface connected from the PC. The alternative method is a remote unicast connection to the VBrick 3200. The client viewer for the VBrick 3200 is a proprietary application that must be installed on a Windows client. The proprietary viewer is a plug-in for the Windows Media Player (WMP). There is a free multicast multimedia stream player by VBrick that requires user registration making it difficult to implement. The player is a plug-in for Windows Media Player and launched from a web browser.

Metric \ Server		VBrick 3200	Score
Streaming Capacity (#)		1	- / (-1)
Bandwidth (Mbps)		3.5 Mbps	+ / (0)
Quality of Audio and Video	- Frames per Second	30	+ / (1)
	- Resolution Size	½ NTSC	+ / (2)
	- Resolution Size	MPEG1	+ / (3)
	- Audio Sampling Rate	Stereo	+ / (4)
Maintainability	- Scalability	No	- / (3)
	- Cost	Affordable	+ / (4)
	- Maintenance	Proprietary	- / (3)
Versatility		Low Versatility Low Maintenance	- / (2)
Security		No	- / (1)
Final Score			1

Figure 11.   VBrick 3200 Metrics Matrix.

4.   General Purpose Server

The following is the metrics statistics for the General Purpose or Generic Server. The streaming capacity of the General Purpose Server is based on the network capacity of the General Purpose Server network bandwidth available. The bandwidth required for streaming the multimedia content to the multicasting address varies from 500 Kbps to 3.5 Mbps. For streaming recorded content, VBrick offers a free application program called StreamPump allowing a Windows-based computer to stream MPEG1 content to a multicasting address.   The number of streams is approximately 50-55 multicast multimedia streams. The bandwidth required to stream the recorded MPEG1 content is 1.5 Mbps. A different program that runs on a Windows2000 Server-based computer is Darwin Streaming Server, free from Apple. Darwin Streaming Server is the same as QuickTime Streaming Server recompiled to run on Windows2000, and therefore has the

same capabilities and remarks as listed in the Apple XServer mentioned previously. The quality of video offers the range of 2 to 30 frames per second. The resolution size is restricted to 320x240 or less. The CODEC format is either MPEG1 or MPEG4, depending on the application. The VBrick StreamPump CODEC is MPEG1; the Apple Darwin Streaming Server CODEC is MPEG4. The audio sampling rates are from 8 KHz to 48 KHz and allows for stereo recording. The General Purpose Server provides scalability for adding more memory, hard drive storage, and network cards for greater network bandwidth. The cost of the server is in the budget category; however, high-end computers are more expensive, moving the General Purpose Server to Affordable or Expensive, but providing more services. The maintenance is open source and responsibility of the IT Department. The General Purpose Server offers high versatility providing multiple services such as ftp, e-mail, and web pages; but reduces the amount of streams provided by consumed network capacity of other services. The server offers high versatility and requires a high amount of administration to offer the multimedia streaming content. There are no security features that provide encryption to the multimedia streaming content. The server provides other security features expected from an application server.

Here are some additional notes on the General Purpose Server that need to be mentioned. The Darwin Streaming Server has problems as mentioned in the Apple XServer metrics review. Apple maintains and monitors a mailing list for Darwin Streaming Server where individuals post questions for additional help, usually answered by someone who had a similar problem or by an Apple Representative. Darwin Streaming Server is available for RedHat Linux. The General Purpose Server was not reconfigured for RedHat Linux for testing. The application provides the same features of the Apple XServer version; the underlying operating system is different. Using RedHat Linux keeps the overall costs in the Budget category because the operating system is available for free. This configuration does not provide live multicast multimedia streaming. QuickTime Broadcaster is not available for RedHat Linux or Windows.

The free version of StreamPump by VBrick is a file server that streams multimedia content to a multicasting address. There is no support provided by VBrick unless purchased. The proprietary version available from VBrick offers more capabilities

to the multicast stream. VBrick was contacted on multiple occasions for possible evaluation of additional equipment and configurations; however, VBrick never responded to any of the attempted contacts.

Metric \ Server		General Purpose Server	Score
Streaming Capacity (#)		Limited by Network Bandwidth	+ / (1)
Bandwidth (Mbps)		1.5 Mbps	+ / (2)
Quality of Audio and Video	- Frames per Second	30	+ / (3)
	- Resolution Size	½ NTSC	- / (2)
	- Resolution Size	MPEG1, MPEG4	+ / (3)
	- Audio Sampling Rate	Stereo	+ / (4)
Maintainability	- Scalability	Yes	+ / (5)
	- Cost	Budget	+ / (6)
	- Maintenance	COTS	- / (5)
Versatility		High Versatility High Maintenance	- / (4)
Security		No	- / (3)
Final Score			3

Figure 12.   General Purpose Server Metrics Matrix.

Table 3 provides a side-by-side comparison of the four Streaming Multimedia Multicasting Servers.

Server	Streaming Capacity	Bandwidth (Mbps)	Quality of Audio / Video				Maintainability			Versatility	Security
			FPS	Resolution	CODEC	Audio	Scalability	Cost	Maintenance		
Apple XServer	Live -1 60-65	0.8 – 3.5	30	½ NTSC	MPEG4	Stereo	Yes	Afford.	Proprietary	High Vers. High Maint.	No
Video Furnace Server	1	0.5 – 7	30	FULL NTSC	MPEG1	Stereo	Yes	Exp.	Proprietary	High Vers. Low Maint.	Yes
VBrick 3200	1	1 – 3.5	30	½ NTSC	MPEG1	Stereo	No	Afford.	Proprietary	Low Vers. Low Maint.	No
General Purpose Server	50-55	1.5	30	½ NTSC	MPEG1 MPEG4	Stereo	Yes	Budget	COTS	High Vers. High Maint.	No

Table 3.   Server Metrics Comparison

B.   CONCLUSION

The thesis provided the necessary background, discussed the comparison of Streaming Multimedia Multicasting Servers, the metrics that should be used for comparison, and the methodology for comparing different systems available. Though the thesis shows several examples of the methodology used to evaluate the Streaming Multimedia Multicasting Servers, the end result will come down to what the IT Department, administrators, and technicians are comfortable with installing and maintaining. The thesis is just a recommendation and should not be the final decision. There are difficulties when trying to compare various types of hardware and software; hopefully the thesis provided the necessary clarification for a starting point in implementing a multicast server.

The thesis presents a brief background of networking, multimedia, multicasting, and streaming. This provides a starting point when deciding on a multicasting server for Distance Learning. The metrics defined gives a baseline of measurements to look for when comparing various server platforms. The network must be multicast enabled to reap the full benefits.

The Streaming Multimedia Multicast Server evaluated with the best performance to price ratio is the Video Furnace server for streaming multimedia content using the multicasting protocol. Although the Video Furnace multicast server is expensive, the implementation, administration, and support will prove beneficial and result in a positive return on the investment (ROI). The client computer does not require an application program or plug-in; the Video Furnace server provides "on the fly" the small application required to view the multicast multimedia stream for the appropriate operating system of the client. Once the client application is closed, no footprint of data left on the client computer. The client operating system is not limited to one provider, but covers Windows, Linux, and Apple versions.

The other Streaming Multimedia Multicast Serves evaluated in the thesis offer great alternatives at a lower price. The servers may or may not have all the benefits, features, or services that are desired or required. The VBrick 3200 does not offer additional services or upgradeability, but offers ease of use. The Apple XServer and

General Purpose Server are more administrative intensive, but offer the flexibility to provide other services in one computer. One of the features to weigh is the encryption of the multicasting multimedia stream. If encryption is not needed, one of the other multicast multimedia servers may prove to be a better selection. The last area for consideration, though not measured, is the client application used with the multicast multimedia server. With a proprietary application that requires installation on every client, that places added demands on the IT Department staff to ensure the client computers have it installed.

C.   RECOMMENDATIONS

The recommendation to the Naval Postgraduate School Code 05 IT Department is to start with several streaming multimedia multicast servers for distributed learning on the main campus using the bandwidth offered by Internet2. Phase in additional multicasting multimedia servers as the demands for distance learning increases and proves beneficial. The first implementation should be to install servers by Video Furnace in the Distance Learning classrooms and one in the Network Operations Center. The server offered by Video Furnace provides ease of implementation and maintenance for a multicasting server. The Distance Learning classrooms in Root Hall would provide a great installation base for the multicasting servers since the classrooms are established Distance Learning classrooms with cameras. The camera audio/video can be split keeping the original setup and the other routed to the Video Furnace multicasting server. These classrooms are overcrowded during lectures with students going to a remote classroom to receive the lecture. With the implementation of the Video Furnace multicasting server, the students that can not find seats in the classroom can go to any computer lab and view the lecture or the remote classroom as established. A separate server can be installed in the Network Operations Center (NOC) for playback of recorded lectures and mandatory training videos that require periodic viewing of all military staff and students. The VF OnGuide is easily updated to reflect the lectures being streamed out via multicasting. When the NOC Video Furnace multicasting server is not being used for streaming a training video or lecture, the server can be used to stream CNN via multicasting. With CNN being one of the authorized channels for watching during working hours, the staff and students will be able to keep current with US and world

events. The staff members and students will be able to view CNN without placing a strain on the campus network due to the benefits multicasting. The Video Furnace server has the capability to redirect all channels listed in the VF OnGuide to one multicast multimedia stream for viewing.

If the Video Furnace Multicast Server is implemented on the expansion Internet2 .edu connection of NPS, their may be additional multicast channels available from other institutions. California State University of Monterey Bay has plans to implement a Video Furnace Multicast Server as part of their network expansion onto the California Research and Education Network (CalREN) and Internet2 connections. The Distance Learning Center of NPS can be expanded to incorporate the lectures provided by other educational institutions.

D.   FUTURE WORK

There are several areas of future work that can be expanded upon from this thesis. The different areas are implementation of the multimedia multicasting server, researching multicasting over IPv6, researching multicasting over wireless, and researching the new MPEG standards that are being released.

Implementing the Streaming Multimedia Multicasting Server to provide the multimedia content to various classrooms and computer labs throughout the Naval Postgraduate School campus will use a small percentage of the new bandwidth provided by the Internet2 expansion connection. This will require detailed planning with the IT Department and the NPS Streaming Media Services Committee.

Research and implement a Streaming Multimedia Multicasting Server using RedHat Linux and Darwin Streaming Server. Research and implement a Streaming Multimedia Multicasting Server with Windows 2003 Server and Darwin Streaming Server. Conduct a comparison of the systems for performance between streaming multicast while multiple services are running.

Further research will be required to implement IP multicasting on the NPS network using IPv6. The Department of Defense will be required to have all networks IPv6 compliant. IPv6 is the Internet Protocol version designed to phase out IPv4 and the problems with network addressing, subnet masking, and security.

There are new MPEG standards that are being released that allow for different streaming bandwidth requirements that allow for dial-up connections and wireless handheld devices to receive multimedia content. The latest MPEG standards are MPEG7 and MPEG21. Research could be done on the new versions of the MPEG standards and how the new MPEG standard will be implemented using the multicasting protocol and IPv6.

With wireless networking becoming prevalent and a standard part of the IT Department and networking, research needs to be conducted on ways to implement multicasting over wireless networking. Multicasting provides bandwidth benefits to the wired networking community; the proposal and implementation of streaming multimedia via multicasting and wireless networking need research.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

[BRAI31]   Brain, Marshall. "How Television Works." [http://entertainment.howstuffworks.com/tv.htm]. July 2003.

[BRAI32]   Brain, Marshall. "How Digital Television Works." [http://entertainment.howstuffworks.com/dtv.htm]. July 2003.

[BROW03]   Brown, Gary. "How HDTV Works." [http://entertainment.howstuffworks.com/hdtv.htm]. July 2003.

[CASE03]   Case, Loyd. "What It Takes To Test A Product." [http://www.extremetech.com/article2/0,3973,771005,00.asp]. May 2003.

[FLOU99]   Flournoy, M. A., *Benefits Of Multicasting Applications Within The United States Marine Corps*. Master's Thesis, Naval Postgraduate School, Monterey, California, June 1999.

[HOST03]   *Hostcentric Streaming Guide: The Basics*. Hostcentric, Inc. 2003.

[KO01]   Susan Ko and Steve Rossen, *Teaching Online A Practical Guide*, Houghton Mifflin, 2001.

[KURO01]   James F. Kurose and Keith W. Ross, *Computer Networking A Top-Down Approach Featuring the Internet*, Addison-Wesley, 2001.

[MACK02]   Steve Mack, *Streaming Media Bible*, Hungry-Minds, 2002.

[MICR03]   Microsoft Encarta Dictionary, [http://encarta.msn.com/encnet/features/dictionary/DictionaryResults.aspx?search=multimedia]. May 2003.

[MILL99]   C. Kenneth Miller, *Multicast Networking and Applications*, Addison-Wesley, 1999.

[QUIC02]   *QuickTime Streaming Server, Darwin Streaming Server: Administrator's Guide*. Apple Computer, Inc. 2002.

[RFC3170]   Network Working Group Request For Comments 3170, *IP Multicast Applications: Challenges and Solutions*, by B. Quinn and K. Almeroth, The Internet Society, 2001.

[VIDE03]   *Video Furnace User's Guide 2.1.* Video Furnace, Inc. 2003.

[WILL00]   Beau Williamson, *Developing IP Multicast Networks Vol I*, Cisco Press, 2000.

[WITT99]   Ralph Wittman and Martina Zitterbart, *Multicast Communication Protocols and Applications*, Morgan Kaufmann, 1999.

[ZACK01]   Craig Zacker, *Networking: The Complete Reference*, Osborne/McGraw-Hill, 2001.

INITIAL DISTRIBUTION LIST

1.   Defense Technical Information Center
        Ft. Belvoir, Virginia

2.   Dudley Knox Library
        Naval Postgraduate School
        Monterey, California

3.   Geoffrey Xie
        Naval Postgraduate School
        Monterey, California

4.   John Gibson
        Naval Postgraduate School
        Monterey, California

5.   Scott Cote
        Naval Postgraduate School
        Monterey, California

6.   Brian Steckler
        Video Furnace, Inc.
        Libertyville, IL

7.   Lonna Sherwin
        Naval Postgraduate School
        Monterey, California

8.   Bob Fimiani
        Apple Federal
        Reston, Virginia

9.   Howard Weinzimmer
        Video Furnace, Inc.
        Libertyville, Illinois

10.   Rick Mavrogeanes
        VBrick Systems, Inc.
        Wallingford, Connecticut